

Measurement-Driven Hysteresis-Aware Friction Identification for Harmonic-Drive Robot Joints Using Symbolic Regression

Zhuoheng Li, Lijia Chen, and Lei Wang, *Member, IEEE*

Abstract—Accurate residual-torque measurement and friction identification are essential for high-performance robot joint modeling, yet harmonic-drive transmissions exhibit reversal-induced hysteresis peaks that are difficult to characterize from current-based torque data. Traditional parametric friction models are structurally rigid, whereas neural-network models often lack interpretability and extrapolate poorly outside the measured operating region. This paper proposes a measurement-driven hysteresis-aware symbolic identification framework for harmonic-drive robot joints. First, rigid-body dynamics are identified and decoupled from current-based torque measurements to obtain residual torque, so that the symbolic model is fitted to the measured non-rigid residual rather than to raw motor-current data. Second, a physics-informed hysteresis state and its derivative are introduced to augment velocity with reversal-memory information. Third, the ParFam symbolic regression algorithm is used to identify compact explicit residual-torque equations from the augmented data. Experiments on a UFACTORY-850 manipulator show that the proposed ParFam-H model achieves the lowest offline residual-torque prediction RMSE among LuGre, Stribeck, RBFNN, RBFNN-H, and velocity-only symbolic baselines. In ± 0.1 s reversal windows, ParFam-H reduces RMSE by up to 45.5% compared with a velocity-only RBFNN, demonstrating improved characterization of transient reversal peaks. Variable-frequency Chirp validation further indicates stable out-of-distribution generalization from 0.1 to 1.0 Hz. Finally, because the learned model is an explicit algebraic equation with only a bounded hysteresis-state update, closed-loop experiments verify its real-time deployability, reducing trajectory-tracking RMSE by 49.3% when used as a feedforward compensation term.

Index Terms—Dynamic identification, friction modeling, harmonic drive, hysteresis, instrumentation and measurement, robot manipulators, symbolic regression.

I. INTRODUCTION

Accurate measurement-driven dynamic identification is a prerequisite for high-fidelity robot modeling, condition monitoring, and high-performance motion control. Although rigid-body robot dynamics can be expressed in a linear regressor form with respect to inertial parameters [1], [2], the torque measured from industrial robot joints contains not only rigid-body effects but also friction, transmission nonlinearities, current-estimation noise, and numerical differentiation errors. Therefore, the reliability of the final model depends critically on how the residual torque is measured, decoupled, and characterized.

This work was supported by the National Natural Science Foundation of China under Grant 62573014.

The authors are with the School of Automation Science and Electrical Engineering, Beihang University, Beijing 100191, China.

Corresponding author: Lei Wang (e-mail: lwang@buaa.edu.cn).

This problem becomes more challenging for harmonic-drive robot joints. Flexspline deformation and gear meshing introduce nonlinear stiffness, position-dependent torque ripple, and hysteretic transmission behavior [3]–[5]. Around velocity reversals, the measured residual torque often contains sharp static-friction peaks and delayed relaxation effects. These reversal regions are precisely where low-speed positioning, force interaction, and dynamic trajectory tracking are most sensitive, but they are also the regions where current-based torque measurement and rigid-body decoupling are most vulnerable to contamination.

Existing friction identification methods commonly use pre-defined parametric structures such as Coulomb/Stribeck or LuGre-type models [10], [11]. These models are interpretable, but their fixed forms cannot easily absorb the coupled nonlinearities of harmonic-drive joints. Data-driven models, including neural networks [13]–[16], provide stronger approximation ability, but they are often black-box predictors and may generate non-physical behavior under unmeasured operating frequencies. Recent continuous dynamic modeling studies also show that friction characterization remains a central limitation in high-precision mechatronic systems [9].

Symbolic Regression (SR) offers a promising alternative because it identifies explicit equations directly from measurement data [17]–[19]. However, most SR-based robotic friction studies still treat friction as a static mapping of velocity [20], which cannot represent the multi-valued residual-torque loops induced by hysteresis. In addition, GP-based SR is prone to inefficient search and equation bloating [21]. These limitations motivate a measurement-driven identification framework that explicitly targets residual-torque behavior near velocity reversals while preserving an interpretable analytical structure.

To address these challenges, this paper proposes HESR, a hysteresis-enhanced symbolic identification framework for harmonic-drive robot joints. The method first extracts residual torque from current-based joint torque measurements after rigid-body dynamic decoupling. It then augments the measured velocity with a physics-informed, measurement-derived hysteresis state and its derivative, transforming path-dependent residual-torque behavior into an augmented Markovian feature space. Finally, the gradient-based ParFam algorithm [22] identifies compact symbolic residual-torque equations from the measured data. Experiments on the first three joints of a UFACTORY-850 manipulator show that the resulting ParFam-H model improves residual-torque prediction, reversal-region characterization, cross-frequency generalization, and closed-

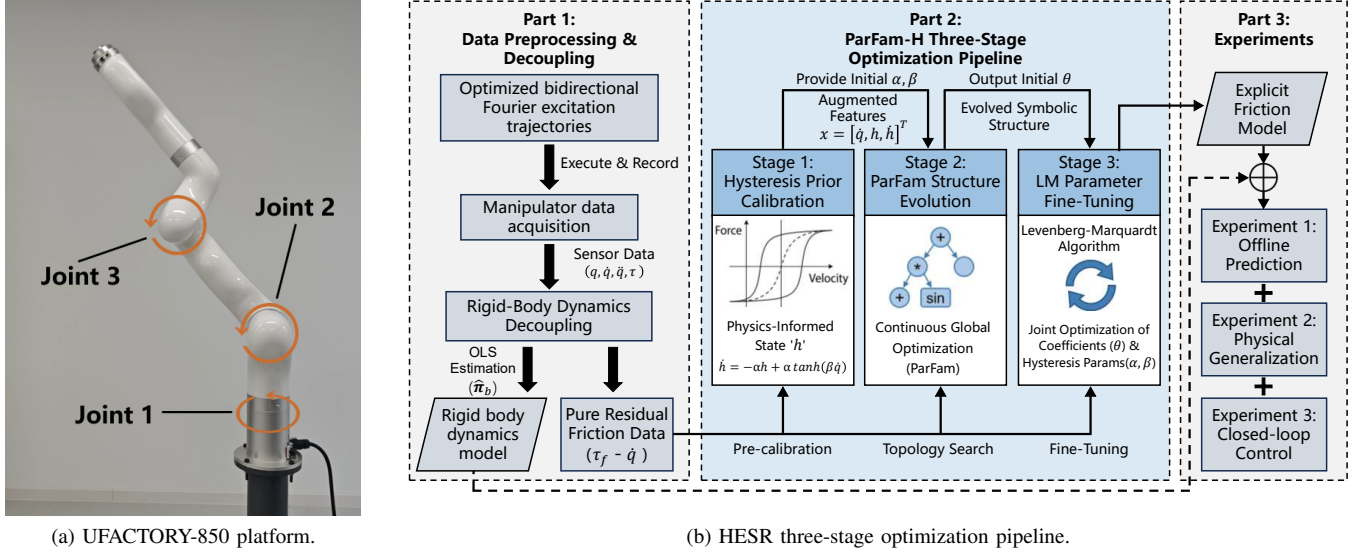


Fig. 1. The overall framework of the proposed HESR methodology. (a) The 6-DOF robotic manipulator (UFACTORY-850) used for data collection; the first three actuated joints are analyzed in this study. (b) The detailed measurement and identification pipeline, comprising three parts: Part 1 handles current-based torque preprocessing and rigid-body dynamic decoupling; Part 2 executes the ParFam-H optimization (including hysteresis prior calibration, structure evolution, and parameter fine-tuning) to identify explicit residual-torque equations; and Part 3 conducts experimental validation.

loop application performance.

The main contributions of this paper are summarized as follows:

- A measurement-driven residual-torque identification pipeline is developed for harmonic-drive robot joints. Current-based torque measurement is first decoupled from rigid-body dynamics, and reversal-region screening is applied only during inertial identification to prevent hysteresis peaks from contaminating the rigid-body estimate.
- A hysteresis-state augmented symbolic-regression formulation is proposed. By reconstructing h and \dot{h} deterministically from the measured velocity \dot{q} , the method converts multi-valued reversal behavior into an augmented feature space suitable for explicit equation discovery.
- A ParFam-based residual-torque model is identified as a closed-form algebraic equation, providing an interpretable and real-time deployable alternative to black-box neural friction predictors.
- A reversal-focused validation protocol is provided, including offline residual-torque prediction, ± 0.1 s reversal-window error analysis, threshold and filtering sensitivity checks, variable-frequency Chirp testing, computation-time analysis, and closed-loop application experiments.

The complete measurement and identification pipeline is shown in Fig. 1. The remainder of this paper is organized as follows. Section II reviews related work on robot dynamic identification, friction modeling, and symbolic regression. Section III presents the residual-torque measurement and rigid-body decoupling procedure. Section IV introduces the hysteresis-aware symbolic identification method. Section V validates the identified model through data traceability checks,

prediction, reversal-window characterization, cross-frequency testing, and closed-loop application experiments. Section VI concludes the paper.

II. RELATED WORK

A. Robot Dynamic Identification and Residual Torque Measurement

Robot dynamic identification commonly starts from Newton-Euler or Euler-Lagrange dynamics and rewrites the inverse dynamics as a linear regressor with respect to inertial parameters [1], [2]. Since the standard inertial parameter vector contains redundant combinations, QR decomposition or Singular Value Decomposition (SVD) is typically used to extract a base parameter set before least-squares estimation. To improve numerical conditioning, finite-Fourier excitation trajectories are widely adopted, often by minimizing the condition number of the observation matrix.

For industrial robots, however, the measured joint torque is not a pure rigid-body quantity. It is usually obtained from motor current feedback and therefore includes current-to-torque conversion errors, drive-chain nonlinearities, friction, numerical differentiation noise, and filtering effects. As a result, residual torque measurement and rigid-body decoupling become inseparable from dynamic identification. Forward-reverse trajectory averaging is often used to reduce velocity-odd friction components before estimating inertial parameters, but this assumption becomes fragile around velocity reversals where hysteresis breaks odd symmetry.

Recent robot identification studies have incorporated nonlinear friction models, adaptive data acquisition, and optimization-based parameter estimation into this pipeline. For example, Dong *et al.* [6] identified industrial robot dynamics using a nonlinear friction model and an LS-SOS algorithm,

combining friction-parameter estimation, Fourier excitation, base-parameter selection, and global parameter optimization. Zhou *et al.* [7] further introduced a semilinearized dynamic identification approach based on nonlinear friction modeling, while Wang *et al.* [8] investigated adaptive acquisition and least-angle-regression-based dynamic parameter identification for industrial robots. Such methods show that friction-aware identification improves torque-estimation accuracy. Nevertheless, they still rely on pre-defined friction structures and do not explicitly characterize reversal-induced residual-torque peaks from measurement data. This motivates a measurement-driven identification method that can preserve the interpretability of dynamic models while adapting the friction structure to observed harmonic-drive behavior.

B. Dynamic Friction Modeling for Robot Joints

Friction modeling is generally categorized into static and dynamic approaches. Static models, such as Coulomb and Stribeck models [23], assume that friction is a single-valued function of velocity. These models can describe steady-state viscous and low-speed Stribeck effects, but they cannot represent path-dependent residual-torque loops or reversal transients. This limitation is particularly important for harmonic-drive robot joints, where flexspline deformation, meshing error, and torsional compliance introduce hysteresis and position-dependent torque ripple [3]–[5].

Dynamic friction models address memory effects by introducing internal states. The LuGre model [11], for instance, uses a bristle-deflection state z to describe pre-sliding displacement and friction memory. The friction torque τ_f is modeled as:

$$\tau_f = \sigma_0 z + \sigma_1 \dot{z} + \sigma_2 \dot{q} \quad (1)$$

where σ_0 , σ_1 , and σ_2 represent the stiffness coefficient, microscopic damping coefficient, and viscous friction coefficient, respectively. The evolution of the bristle state z is governed by:

$$\dot{z} = \dot{q} - \frac{\sigma_0 |\dot{q}|}{g(\dot{q})} z \quad (2)$$

Here, the function $g(\dot{q})$ captures the Stribeck effect:

$$g(\dot{q}) = F_c + (F_s - F_c) e^{-(\dot{q}/v_s)^2} \quad (3)$$

where F_c is Coulomb friction, F_s is static friction, and v_s is the Stribeck velocity. Friction identification and compensation have also been studied from an instrumentation perspective, where accurate friction-parameter measurement directly affects manipulator compensation performance [12]. While LuGre and related dynamic friction models unify pre-sliding, hysteresis, and steady-state behavior, their fixed mathematical structures can be restrictive for harmonic-drive joints whose residual torque contains coupled transmission ripple, elastic deformation, and static-friction peaks. Therefore, an effective identification method should preserve the memory mechanism needed around reversals while allowing the algebraic friction structure to be discovered from measured residual torque.

C. Data-Driven and Symbolic Identification

Data-driven approaches, especially neural networks, have been widely used for nonlinear friction and robot dynamics identification [13], [24]. Radial Basis Function Neural Networks (RBFNNs) provide efficient static function approximation, while recurrent architectures such as LSTM networks can encode motion history for dynamic friction [14], [25], [26]. Physics-Informed Neural Networks (PINNs) further attempt to embed physical constraints into learning [27], [28]. Despite their approximation capability, these models are often difficult to interpret, may extrapolate poorly under unmeasured frequencies, and can be computationally burdensome for high-rate industrial controllers.

Symbolic Regression (SR) bridges mechanistic and data-driven identification by discovering explicit analytical equations from data [17]–[19]. It has been explored for robot inverse dynamics [29], physical law discovery, and data-efficient adaptive control [30]. In robotic friction modeling, however, existing SR studies often focus on static velocity-to-friction mappings [20], [31]. Such mappings cannot represent the multi-valued residual-torque behavior observed near velocity zero-crossings. Moreover, traditional GP-based SR suffers from slow convergence, sensitivity to noise, and bloated equations [21]. The gradient-based ParFam algorithm [22] offers a more compact continuous-optimization alternative. Building on this idea, the proposed method introduces a physics-informed hysteresis state before symbolic identification, so that ParFam searches for explicit equations in an augmented feature space capable of representing measured reversal memory.

III. RESIDUAL TORQUE MEASUREMENT AND DYNAMIC DECOUPLING

The proposed identification framework starts from measured joint torque and separates the rigid-body contribution from the residual torque. This section describes the current-based torque measurement, rigid-body parameter identification, residual-torque extraction, and reversal-region screening procedure used before symbolic residual-torque identification.

A. Current-Based Torque Measurement

The UFACTORY-850 platform does not provide direct joint torque sensing. Therefore, the measured joint torque τ_{meas} is obtained from motor-current feedback after transmission-ratio and torque-constant conversion:

$$\tau_{meas}(t) = \mathbf{K}_\tau \mathbf{N} \mathbf{i}(t), \quad (4)$$

where $\mathbf{i}(t)$ is the measured motor-current vector, \mathbf{K}_τ is the motor torque-constant matrix, and \mathbf{N} represents the joint transmission ratio. The robot state is sampled at 200 Hz. Joint velocity and acceleration are obtained from the measured position through numerical differentiation and filtering. Consequently, τ_{meas} contains not only rigid-body torque but also friction, current-estimation noise, transmission nonlinearities, and differentiation-induced errors. The objective of the following steps is to extract a reliable residual-torque signal from this measured signal.

B. Rigid-Body Dynamic Identification

For the spatial 3-DOF manipulator configuration investigated in this study, the rigid-body dynamics are described by the standard Euler-Lagrange equation:

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q}) = \boldsymbol{\tau}_{rb} \quad (5)$$

where $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}} \in \mathbb{R}^3$ are the joint position, velocity, and acceleration vectors, respectively. $\mathbf{M}(\mathbf{q})$ is the inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ contains Coriolis and centrifugal terms, $\mathbf{G}(\mathbf{q})$ represents gravity, and $\boldsymbol{\tau}_{rb}$ is the rigid-body torque.

Exploiting the linear property of robot dynamics with respect to inertial parameters, (5) can be rewritten in a regressor form:

$$\mathbf{Y}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})\boldsymbol{\pi} = \boldsymbol{\tau}_{rb} \quad (6)$$

where $\boldsymbol{\pi} \in \mathbb{R}^{30}$ is the standard inertial parameter vector (including mass, center of mass, and inertia tensor components for each link), and \mathbf{Y} is the regression matrix.

Due to kinematic constraints, linear dependencies exist within $\boldsymbol{\pi}$, making direct Least Squares estimation ill-posed. Therefore, a minimum parameter set $\boldsymbol{\pi}_b$ must be extracted. Following the standard base-parameter identification procedure for robot dynamics [2], we construct an extended observation matrix using randomly generated valid joint states and apply numerical QR decomposition with column pivoting to remove linearly dependent columns. For our spatial 3-DOF configuration, the original 30 standard parameters are reduced to a base set of 15 parameters. This number is not assumed a priori; it is determined by the numerical rank of the extended observation matrix under the adopted kinematic chain and is consistent with the 15-column minimal regressor used in all subsequent torque reconstruction and residual-torque extraction experiments. The minimal dynamics model is thus formulated as:

$$\mathbf{Y}_b(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})\boldsymbol{\pi}_b = \boldsymbol{\tau}_{rb} \quad (7)$$

C. Excitation Trajectory and Rigid-Body Decoupling

To enhance the robustness of parameter identification against measurement noise, finite Fourier series are adopted to construct periodic excitation trajectories. The full observation matrix \mathbf{W}_b is constructed by vertically stacking the regressor \mathbf{Y}_b calculated at each sampling instant along the trajectory. Consequently, the fundamental frequency and Fourier coefficients are optimized by minimizing the condition number of \mathbf{W}_b .

During the estimation phase, a forward-reverse symmetric motion strategy is implemented to decouple friction from the rigid-body dynamics. The manipulator is controlled to execute a forward trajectory $\mathbf{q}^+(t)$ and its time-reversed counterpart $\mathbf{q}^-(t)$. Theoretically, the odd-symmetric nature of friction with respect to velocity allows it to be canceled out by averaging the torques from both directions.

By retaining only the data segments with well-preserved friction symmetry, the high-confidence rigid-body torque $\bar{\boldsymbol{\tau}}_{rb}$ is obtained. The optimal estimate of the base parameters $\hat{\boldsymbol{\pi}}_b$ is computed using Ordinary Least Squares (OLS):

$$\hat{\boldsymbol{\pi}}_b = (\mathbf{W}_b^\top \mathbf{W}_b)^{-1} \mathbf{W}_b^\top \bar{\boldsymbol{\tau}}_{rb} \quad (8)$$

D. Residual-Torque Extraction

Once $\hat{\boldsymbol{\pi}}_b$ is determined, the measured residual torque can be extracted from additional sinusoidal sweep data. Using the identified rigid-body model, the residual torque is defined as:

$$\boldsymbol{\tau}_f = \boldsymbol{\tau}_{meas} - \mathbf{Y}_b(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})\hat{\boldsymbol{\pi}}_b \quad (9)$$

where $\boldsymbol{\tau}_{meas}$ is the current-based joint torque measurement. This residual signal is not merely a fitted numerical error; it is the measured torque component remaining after rigid-body dynamic decoupling. It predominantly contains Coulomb friction, viscous friction, static-friction transients, harmonic-drive hysteresis, and current-estimation noise. Therefore, the subsequent symbolic model is identified directly from measured residual-torque data.

E. Reversal-Region Data Screening

The forward-reverse averaging strategy assumes that friction is approximately odd with respect to velocity. However, this assumption fails near velocity reversals because hysteresis introduces delayed relaxation and asymmetric static-friction peaks. Directly averaging forward and reverse torques in this narrow region can inject hysteresis residuals into the rigid-body estimate rather than cancel them.

To prevent this contamination, data points satisfying $|\dot{q}| < 0.1$ rad/s are excluded only during rigid-body parameter identification. This threshold is selected after checking neighboring values. Repeating the identification with 0.05, 0.10, and 0.15 rad/s produced the same active base-parameter set and changed the reconstructed rigid-body torque by less than the current-derived torque noise level in the retained excitation segments. Importantly, these low-velocity samples are not removed from the friction-identification dataset. They are retained after rigid-body decoupling because they contain the reversal-induced hysteresis loops and static-friction peaks that the proposed symbolic model is designed to characterize.

IV. HYSTERESIS-AWARE SYMBOLIC RESIDUAL-TORQUE IDENTIFICATION

Based on the measured residual torque $\boldsymbol{\tau}_f$ extracted in Section III, this section presents the proposed HESR identification framework. The goal is not only to fit the measured residual torque accurately, but also to identify an explicit residual-torque equation that is interpretable, numerically stable, and suitable for real-time deployment. HESR consists of three stages: 1) *Hysteresis Prior Pre-calibration*, where the shape parameters (α and β) of the physics-informed hysteresis state are initialized from measured residual-torque loops; 2) *Symbolic Structure Evolution*, where ParFam searches for a compact analytical residual-torque equation in an augmented feature space; and 3) *Parameter Fine-Tuning*, where Levenberg-Marquardt (LM) jointly refines the hysteresis parameters and symbolic coefficients.

A. Hysteresis State for Residual-Torque Memory

Most symbolic residual-torque identifiers rely solely on joint velocity \dot{q} as the input, implicitly assuming a static mapping

$\tau_f = f(\dot{q})$. This assumption is inconsistent with the measured residual-torque behavior of harmonic-drive joints, where the same velocity can correspond to different torque values depending on whether the joint is accelerating, decelerating, or crossing zero velocity. Inspired by dynamic friction models such as LuGre [11], which use an internal state to describe bristle deflection, we introduce a compact hysteresis state before symbolic identification.

However, the LuGre model presents two major limitations in practical harmonic drive control: 1) *Structural rigidity*: it is mathematically hardcoded strictly for idealized friction, restricting its capacity to absorb other complex, unmodeled nonlinearities (e.g., flexspline elastic deformation); and 2) *Parameter overload*: its numerous coupled coefficients are notoriously difficult to identify accurately.

Instead of pre-defining a rigid, parameter-heavy physical friction law, we formulate a minimalist hysteresis state h . This state is not an additional sensor measurement; it is a measurement-derived virtual state reconstructed deterministically from the measured joint velocity. It is governed by a first-order nonlinear differential equation with two physically meaningful parameters (α and β), as illustrated in Fig. 2:

$$\dot{h} = -\alpha h + \alpha \tanh(\beta \dot{q}) \quad (10)$$

where $\alpha > 0$ and $\beta > 0$ are tunable physical parameters. In a discrete real-time control system with sampling time Δt , the state is iteratively updated via first-order Euler discretization:

$$\begin{cases} h_k = (1 - \alpha \Delta t) h_{k-1} + \alpha \Delta t \tanh(\beta \dot{q}_k) \\ \dot{h}_k = -\alpha h_k + \alpha \tanh(\beta \dot{q}_k) \end{cases} \quad (11)$$

where $0 < 1 - \alpha \Delta t < 1$ is maintained to ensure numerical stability and non-oscillatory behavior.

This minimalist design separates the memory prior from the algebraic residual-torque identification. Equation (10) generates the phase-delayed memory coordinate from the measured velocity, while ParFam identifies how \dot{q} , h , and \dot{h} combine into the measured residual torque. In this way, the model avoids the parameter burden of LuGre while preserving the state information needed for reversal-region identification. Equivalently, the non-local memory of the friction process is compressed into an augmented Markovian feature vector, allowing a static symbolic equation to represent path-dependent residual-torque behavior.

Furthermore, the structural design of (10) is based on the microscopic physical interactions of harmonic drives. From a tribological perspective, the gear meshing contact surfaces can be modeled as interacting elastic asperities (bristles). First, the transition term $\tanh(\beta \dot{q})$ replaces the discontinuous sign function with a smooth transition around zero velocity. A larger β increases the local slope near $\dot{q} = 0$, reflecting the abrupt asperity break-away during the Stribeck transition. Second, the decay term $-\alpha h$ governs the micro-elastic relaxation process and can be interpreted as a first-order low-pass memory with bandwidth controlled by α . When the joint velocity reverses and crosses zero, the microscopically deformed bristles do not instantly snap back to a neutral state. Instead, they release their stored elastic strain energy exponentially. The parameter

α therefore acts as the inverse of this relaxation time constant. By mathematically preserving this residual micro-deformation state, $h_k \approx (1 - \alpha \Delta t) h_{k-1}$, the variable h physically reconstructs the path-dependent hysteresis loop without the structural complexity of traditional models. The derivative \dot{h} is additionally retained because it indicates whether the residual-torque memory state is undergoing rapid transition, which is especially informative near reversal-induced peaks.

B. Boundedness of the Hysteresis Feature

The measurement-derived hysteresis state is also bounded under any bounded velocity input. Since $|\tanh(\beta \dot{q})| \leq 1$, the continuous dynamics in (10) drive h toward the interval $[-1, 1]$. Specifically, for the Lyapunov candidate $V(h) = h^2/2$,

$$\dot{V} = h\dot{h} = -\alpha h^2 + \alpha h \tanh(\beta \dot{q}) \leq \alpha |h|(1 - |h|). \quad (12)$$

Thus, when $|h| > 1$, $\dot{V} < 0$, and the state is driven back toward the compact invariant set $|h| \leq 1$. This boundedness is important for measurement-driven identification because it prevents the reconstructed feature from unboundedly amplifying velocity noise during long trajectories. Both h and its derivative \dot{h} are therefore used as stable memory features for residual-torque identification.




Fig. 2. Composite hysteresis analysis of the proposed residual-torque memory state governed by Equation (10). (a) The shape factor β dictates the slope of the \tanh function near zero velocity. (b) The time constant α determines the decay rate of the state variable h , reflecting the memory depth and the width of the hysteresis loop.

C. Three-Stage Symbolic Identification Strategy

Residual-torque identification using SR involves both function topology search and parameter identification. To ensure physical consistency, compactness, and robust numerical fitting, HESR uses a physics-guided three-stage identification framework.

Stage 1: Hysteresis Prior Pre-calibration. Before symbolic search, the physical parameters α and β are pre-calibrated from measured $\tau_f - \dot{q}$ residual-torque loops. For each joint, candidate pairs (α, β) are sampled from bounded positive intervals and the corresponding state h is generated by (11). The initial pair is selected by minimizing the mismatch between the measured hysteresis loop and the generated loop descriptor:

$$J_h(\alpha, \beta) = w_a |\Delta \tau_f - \Delta h| + w_s |s_{\tau,0} - s_{h,0}|, \quad (13)$$

where $\Delta\tau_f$ and Δh denote normalized loop widths around velocity reversal, $s_{\tau,0}$ and $s_{h,0}$ denote normalized slopes near $\dot{q} = 0$, and w_a, w_s are fixed weighting factors. This step does not determine the final residual-torque equation; it only initializes a physically meaningful memory coordinate for the subsequent symbolic identification.

Stage 2: ParFam-based Symbolic Structure Evolution. Unlike traditional GP, which searches discrete symbolic spaces and is often prone to code bloat and poor constant optimization, ParFam transforms SR into a continuous global optimization problem over a parametric family of functions:

$$f_{\theta}(\mathbf{x}) = Q_{k+1}(\mathbf{x}, g_1(Q_1(\mathbf{x})), \dots, g_k(Q_k(\mathbf{x}))) \quad (14)$$

where $\mathbf{x} = [\dot{q}, h, \dot{h}]^T$ is the augmented input vector and θ denotes the learnable algebraic coefficients. The operator library is fixed for all joints as $\mathcal{O} = \{+, -, \times, /, \sin, \cos, \exp, \tanh\}$, while ParFam dynamically selects active terms by driving unnecessary coefficients to zero. This setting gives all joints the same search space; the final equations in Table VII are therefore selected from measured data rather than manually assigned. The identification objective is:

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N (\tau_{f,i} - f_{\theta}(\mathbf{x}_i))^2 + \gamma \|\theta\|_1 \quad (15)$$

The L_1 regularization prunes redundant terms by forcing unnecessary coefficients toward zero. We globally optimize θ via Basin-Hopping and L-BFGS-B local search. Basin-Hopping provides random perturbations to escape poor local minima, while L-BFGS-B performs bounded local refinement of continuous coefficients. After the global search, coefficients whose magnitudes fall below a fixed pruning threshold are set to zero, and the remaining active expression is refined by a final BFGS step. To prevent non-physical bloating, polynomial degrees are restricted to 1 and cross-unit operations are prohibited. Candidate divisions are protected by lower-bounding denominators during optimization to avoid singular expressions. This implementation converts the discrete symbolic-search problem into a continuous sparse identification problem and yields concise residual-torque equations suitable for real-time evaluation.

Stage 3: LM-based Joint Parameter Fine-Tuning. Since the pre-calibration in Stage 1 relies on geometric loop descriptors, it inevitably introduces initial bias. After the symbolic structure is fixed in Stage 2, Levenberg-Marquardt (LM) is used to jointly refine the hysteresis parameters (α, β) and the discovered algebraic coefficients θ :

$$\min_{\alpha, \beta, \theta} \sum_{i=1}^N \left(\tau_{f,i} - f_{\theta}(\dot{q}_i, h_i(\alpha, \beta), \dot{h}_i(\alpha, \beta)) \right)^2. \quad (16)$$

This final step removes cascading errors caused by pre-calibration bias and aligns the physical memory coordinate with the measured residual-torque equation.

V. EXPERIMENTAL VALIDATION

Experiments on the first three joints of a UFACTORY-850 cobot (harmonic drives, 200 Hz) validate the proposed

measurement-driven identification framework. The evolved symbolic model is denoted as ParFam-H. The validation is organized around five questions: 1) whether the current-based torque measurement and rigid-body decoupling provide reliable residual-torque data; 2) whether the data processing chain is traceable from raw acquisition to symbolic identification; 3) whether the identified residual-torque equation improves offline prediction over physical, neural, and symbolic baselines; 4) whether the model captures reversal-window peaks that are most sensitive to hysteresis; and 5) whether the identified equation remains useful under frequency extrapolation and closed-loop implementation.

A. Experimental Protocol and Dataset Construction

Table I summarizes the experimental protocol. The same hardware platform, sampling rate, and current-based torque-estimation channel are used throughout the study. The rigid-body identification dataset is collected using an optimized Fourier excitation trajectory and a forward-reverse symmetric execution strategy. The residual-torque identification dataset is then collected using a sinusoidal sweep at 0.25 Hz, from which the identified rigid-body torque is subtracted. Two independent validation datasets are used for offline analysis: an untrained 20 s Fourier trajectory with 4000 samples and a variable-frequency Chirp trajectory from 0.1 to 1.0 Hz. Finally, a 10 s dynamic trajectory tracking experiment is used to evaluate whether the identified residual-torque equation is useful in a real-time application.

TABLE I
EXPERIMENTAL PROTOCOL AND DATASET CONSTRUCTION

Item	Setting
Platform	UFACTORY-850 harmonic-drive cobot, first three joints
Sampling and torque source	200 Hz; joint torque estimated from motor-current feedback
Rigid-body identification	Optimized Fourier excitation with forward-reverse symmetric execution
Reversal screening	Samples satisfying $ \dot{q} < 0.1$ rad/s excluded only from rigid-body identification
Residual-torque training	0.25 Hz sinusoidal sweep after rigid-body torque subtraction
Offline validation	Untrained 20 s Fourier trajectory, 4000 samples
Frequency extrapolation	Chirp trajectory from 0.1 to 1.0 Hz, split at 0.6 Hz
Closed-loop application	10 s dynamic trajectory tracking through velocity-command interface

For every dataset, joint positions and motor-current feedback are recorded synchronously. Joint velocity and acceleration are obtained from the measured position through numerical differentiation and filtering, and the same preprocessing pipeline is used for both training and validation. This separation between the identification trajectory and the validation trajectories is important: ParFam-H is not evaluated on the same motion pattern used to identify the residual-torque equation, and the Chirp trajectory intentionally exceeds the 0.25 Hz training frequency to test cross-frequency generalization.

B. Data-Processing Traceability and Reliability Checks

To make the measurement-driven identification procedure reproducible, the experimental data are organized into three traceable layers, as summarized in Table II. The raw acquisition layer records desired trajectories, measured joint positions, and current-based joint torques from the UFACTORY-850 controller. The dynamic-decoupling layer applies the same filtering and numerical differentiation settings to all records, identifies the rigid-body base parameters from the Fourier excitation data, and subtracts the reconstructed rigid-body torque to obtain residual torque. The symbolic-identification layer then uses only the residual-torque records containing time, velocity, residual torque, h , and \dot{h} to train ParFam-V and ParFam-H. This separation prevents the symbolic model from accessing raw rigid-body torque terms during residual-torque fitting and keeps the reported comparisons tied to the same measured residual signal.

TABLE II
TRACEABLE DATA CHAIN USED FOR MEASUREMENT-DRIVEN IDENTIFICATION

Layer	Main Record	Role in This Study
Raw acquisition	$q_d, \dot{q}, \ddot{q}_d, \ddot{q}_d, \tau_{meas}$	Current-based torque measurement and trajectory execution
Dynamic decoupling	Filtered $q, \dot{q}, \ddot{q}, \tau_{meas}, \tau_{rb}$	Base-parameter identification and residual-torque extraction
Symbolic identification	$t, \dot{q}, \tau_f, h, \dot{h}$	ParFam-V/ParFam-H training and explicit equation discovery

Three reliability checks are used to avoid contradictions between the measurement and identification stages. First, all reported training and validation records are aligned to the 200 Hz controller sampling rate; the 20 s Fourier validation record therefore contains 4000 samples, consistent with Table I. Second, the reversal-screening threshold is applied only when estimating the rigid-body parameters and is not applied to the residual-torque dataset used for ParFam-H, ensuring that the reversal peaks evaluated in the reversal-window analysis are not removed during training. Third, the same differentiation, filtering, and current-to-torque conversion pipeline is applied before comparing LuGre, Stribeck, RBFNN, RBFNN-H, ParFam-V, and ParFam-H. Therefore, the differences in the reported RMSE values arise from the model structures rather than from inconsistent preprocessing.

To verify that the selected reversal-screening threshold is not arbitrary, rigid-body identification was repeated with neighboring thresholds. Such sensitivity checks are consistent with recent robot measurement-system calibration studies, where parameter identifiability and sensitivity analysis are used to assess the reliability of the calibration pipeline [33]. As shown in Table III, all three thresholds preserve the same regressor rank of 15. The 0.05 rad/s setting retains more samples but leaves more low-velocity contamination in the retained set, whereas the 0.15 rad/s setting removes too many samples and increases the condition number. The 0.10 rad/s threshold

provides the best compromise between retained excitation, retained-segment fitting accuracy, and residual-torque stability.

TABLE III
REVERSAL-SCREENING THRESHOLD SENSITIVITY

Th.	Samples	Rank	All	Ret.	$\log_{10} \kappa$
0.05	2611	15	4.2530	3.3154	3.536
0.10	1947	15	4.2850	2.8798	3.627
0.15	1528	15	4.3163	2.5542	3.685

The effect of preprocessing was further checked by re-computing the residual-torque prediction under neighboring moving-average smoothing windows. Table IV reports the mean RMSE over the three joints for both the full trajectory and the reversal samples. The variations are negligible across the tested windows, indicating that the observed reversal behavior is not an artifact of a specific smoothing setting.

TABLE IV
FILTERING SENSITIVITY OF PARFAM-H RESIDUAL-TORQUE PREDICTION

Window	Offline RMSE	Reversal RMSE	Rev. Samples
5	3.8427	4.5842	2050
9	3.8404	4.5803	2050
13	3.8379	4.5750	2051

Because the residual torque is obtained from current-based torque estimation and rigid-body torque subtraction, a residual-level Monte Carlo sensitivity analysis was further conducted to evaluate whether the reported advantage is fragile to measurement uncertainty. The perturbation emulates the combined effect of current-to-torque conversion error, torque-constant and transmission-ratio scale uncertainty, numerical differentiation noise, and rigid-body reconstruction error:

$$\tau_f^{mc} = (1 + s)\tau_f + \epsilon, \quad (17)$$

where s is a zero-mean multiplicative scale perturbation and ϵ is additive Gaussian residual-torque noise. Three uncertainty levels were tested with 200 Monte Carlo trials: mild (1% scale and 1% additive noise), moderate (3%/3%), and severe (5%/5%). As shown in Table V, ParFam-H remains ranked first under all tested uncertainty levels. From the nominal case to the severe case, the mean reversal-window RMSE changes by only 0.24%, indicating that the observed reversal-region advantage is not caused by a fragile realization of the current-based residual-torque measurement pipeline.

TABLE V
MEASUREMENT-UNCERTAINTY SENSITIVITY OF PARFAM-H

Level	Scale/Add.	Offline RMSE	Reversal RMSE
Nominal	0/0	3.8580	4.6655
Mild	1%/1%	3.8577 ± 0.0124	4.6650 ± 0.0138
Moderate	3%/3%	3.8657 ± 0.0416	4.6735 ± 0.0461
Severe	5%/5%	3.8692 ± 0.0611	4.6766 ± 0.0677

The identified ParFam-H equations are also suitable for real-time use. After identification, the online model evaluation requires only the hysteresis-state update in (11) and one explicit algebraic residual-torque equation per joint. No recurrent neural state, iterative optimizer, or symbolic search is

executed during control. Table VI reports the batch-equivalent MATLAB evaluation time of all offline prediction models. All models are far below the 5 ms control period at 200 Hz, and ParFam-H remains among the fastest because it is an explicit algebraic expression.

TABLE VI
BATCH-EQUIVALENT ONLINE EVALUATION TIME

Model	Input	Time / Sample (ms)
LuGre	\dot{q}, t	4.22×10^{-5}
Stribeck	\dot{q}	2.19×10^{-5}
RBFNN	\dot{q}	5.66×10^{-3}
RBFNN-H	\dot{q}, h, \dot{h}	5.79×10^{-3}
ParFam-V	\dot{q}	5.41×10^{-5}
ParFam-H	\dot{q}, h, \dot{h}	1.96×10^{-5}

C. Measurement Platform and Baseline Identification

The experimental platform provides joint positions and motor-current feedback from the embedded controller. Joint torque is estimated from the current feedback and then processed through the dynamic decoupling procedure described in Section III. To isolate the residual-torque component, rigid-body dynamics were first identified using an optimized Fourier trajectory and a forward-reverse symmetric motion strategy. Data within the velocity reversal region ($|\dot{q}| < 0.1$ rad/s) were discarded only for rigid-body identification to avoid contaminating the inertial estimate with hysteresis-dominated residuals. QR decomposition yielded 15 minimum base inertial parameters, which accurately track the macroscopic rigid-body torque (Fig. 3).




Fig. 3. Current-based measured torque and identified rigid-body torque under the Fourier excitation trajectory. The close agreement outside reversal-screened regions supports the reliability of the residual-torque extraction step.

Subsequently, a sinusoidal sweep (frequency 0.25 Hz, amplitude 1.2 rad/s) was executed to collect residual-torque data. By subtracting the identified rigid-body torque from the current-based measured torque, the residual torque τ_f was extracted for each joint. These residual measurements contain the reversal-induced peaks and hysteresis loops that are difficult to observe from velocity alone. Finally, ParFam-H, using the augmented features $\mathbf{x} = [\dot{q}, h, \dot{h}]^T$, was applied to identify an explicit residual-torque equation. The specific analytical formulas discovered for each joint, together with the co-optimized hysteresis parameters (α, β), are listed in Table VII.

For a fair comparison, all baseline models were trained on the same residual-torque measurements. The LuGre and Stribeck parameters were optimized by a Genetic Algorithm

(GA) using the same mean-square residual-torque prediction objective:

$$J_{GA}(\boldsymbol{\eta}) = \frac{1}{N} \sum_{i=1}^N (\tau_{f,i} - \hat{\tau}_{f,i}(\boldsymbol{\eta}))^2, \quad (18)$$

where $\boldsymbol{\eta}$ denotes the parameter vector of the corresponding physical model. Parameter bounds were chosen to enforce positive stiffness, damping, Coulomb friction, static friction, and Stribeck velocity. The GA result was further refined by local least-squares optimization, giving the physical baselines a best-case identification setup rather than weak manual tuning.

TABLE VII
EXPLICIT ANALYTICAL RESIDUAL-TORQUE EQUATIONS IDENTIFIED BY PARFAM-H

Residual-Torque Model Equation	
Joint 1:	$\tau_{f1} = 16.237h_1 - 0.163\dot{h}_1 - 2.128 + 37.715 \tanh(\theta_1)$
	where $\theta_1 = 0.193\dot{q}_1 - 0.606h_1 + 0.064\dot{h}_1 + 0.057$ $(\alpha_1, \beta_1) = (11.37, 256.84)$
Joint 2:	$\tau_{f2} = 3.710\dot{q}_2 + 2.936h_2 + 1.220\dot{h}_2 - 0.279 + 8.290 \cos(\theta_2)$
	where $\theta_2 = -0.865\dot{q}_2 + 1.531h_2 + 0.351\dot{h}_2 + 1.485$ $(\alpha_2, \beta_2) = (19.62, 143.91)$
Joint 3:	$\tau_{f3} = 0.051\dot{q}_3 + 6.614h_3 + 0.238\dot{h}_3 - 0.280 + 6.527 \tanh(\theta_3)$
	where $\theta_3 = 1.014\dot{q}_3 - 2.130h_3 + 0.115\dot{h}_3 + 0.167$ $(\alpha_3, \beta_3) = (8.45, 278.23)$

The identified structures also provide physical clues. The dominant h and \dot{h} terms in all joints represent residual-torque memory and relaxation around velocity reversal. The $\tanh(\theta)$ terms in Joints 1 and 3 behave like smooth saturation functions, consistent with bristle break-away and Stribeck-like transitions. The $\cos(\theta_2)$ term in Joint 2 indicates a periodic component mixed with velocity and hysteresis states; this is consistent with the stronger gravity loading and position-dependent transmission ripple of the second joint in the investigated configuration.

D. Offline Residual-Torque Prediction

To evaluate the identification accuracy of ParFam-H on complex residual-torque dynamics, an offline validation was conducted using an untrained, dynamically rich Fourier series trajectory (duration: 20 s, 4000 samples). Actual joint angles and current-based torques were recorded, with velocity and acceleration derived numerically. The rigid-body torque was then subtracted to produce the validation residual-torque signal. For visual readability and direct correspondence with Fig. 4, Table VIII reports the representative 0–10 s interval containing multiple velocity reversals.

ParFam-H was quantitatively compared against five baselines: 1) classic LuGre and 2) Stribeck models (parameters optimized via Genetic Algorithm); 3) a standard RBF Neural Network (RBFNN) using only velocity input; 4) an RBF

Neural Network (RBFNN-H), enhanced with the hysteresis state input for fair comparison; and 5) ParFam-V, an ablation symbolic model using only velocity input. The standard RBFNN and RBFNN-H adopt the same RBF architecture and training hyperparameters (basis width, target error, maximum number of hidden neurons, and train/test split); their only difference is the input feature set. These baselines are specifically selected for their balance between approximation capacity and computational efficiency, ensuring strict feasibility for 200 Hz real-time control. The RBFNN/RBFNN-H pair isolates whether hysteresis features alone improve a black-box network, whereas the ParFam-V/ParFam-H pair isolates the same effect inside the symbolic regression family. While higher-capacity architectures like LSTMs or PINNs are available, their significant recursive computational overhead on standard industrial microprocessors often precludes their use in zero-latency feedforward tasks.

As shown in Fig. 4 and Table VIII, ParFam-H achieves the highest residual-torque prediction fidelity across all three joints. While traditional models (LuGre, Stribeck) yield overly smoothed predictions, the standard velocity-only RBFNN tends to smooth out the sharp torque peaks near reversals. RBFNN-H, although given the same hysteresis features as ParFam-H, can produce non-physical oscillations and does not consistently improve over the standard RBFNN. This indicates that directly appending hysteresis features to a black-box neural approximator does not guarantee improved reversal prediction; the benefit comes from combining the hysteresis state with a compact symbolic structure. The ablation comparison with ParFam-V further supports the importance of the hysteresis state h for capturing reversal dynamics.

Quantitatively, ParFam-H consistently yields the lowest RMSE in the representative validation interval. Compared with the standard RBFNN, ParFam-H reduces RMSE by 33.7%, 9.4%, and 21.1% on Joints 1–3, respectively. As an explicit, interpretable residual-torque equation, ParFam-H surpasses fitting-oriented black-box neural networks while preserving compactness and physical readability.

TABLE VIII
OFFLINE RESIDUAL-TORQUE PREDICTION RMSE (N·M) ON THE REPRESENTATIVE 0–10 S INTERVAL

Model	Joint 1	Joint 2	Joint 3
LuGre	5.9402	6.3538	2.2408
Stribeck	5.9966	6.4965	3.2691
RBFNN	6.4174	5.1660	2.4862
RBFNN-H	5.3616	7.7653	2.6685
ParFam-V	5.6837	5.2086	2.6340
ParFam-H	4.2523	4.6786	1.9619

E. Reversal-Window Residual-Torque Characterization

Because hysteresis-induced residual-torque peaks are concentrated around velocity reversals, global RMSE can understate the main failure mode of static or weakly structured models. Therefore, we further evaluate all models only inside a narrow reversal window. For each joint, velocity zero-crossing instants are detected from the validation trajectory, and all

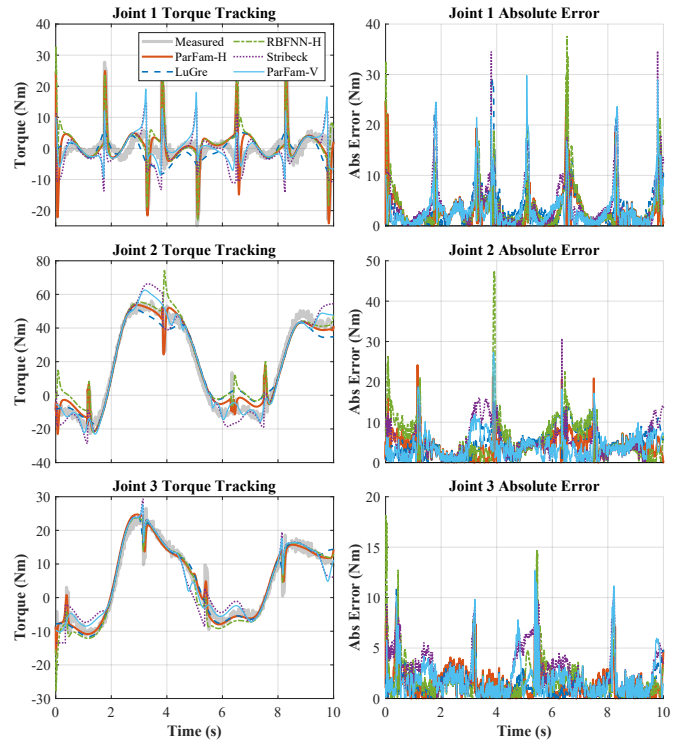


Fig. 4. Residual-torque prediction accuracy and absolute error comparison among different friction models under an untrained, dynamically complex trajectory. For clarity, a representative 0–10 s interval of the 20 s validation dataset is plotted.

samples within ± 0.1 s of each crossing are collected. The resulting metrics are reported in Table IX. This analysis directly targets the measurement region where stick-slip transitions, delayed bristle relaxation, and static-friction peaks dominate.

TABLE IX
RESIDUAL-TORQUE PREDICTION RMSE (N·M) WITHIN ± 0.1 S VELOCITY-REVERSAL WINDOWS

Model	Joint 1	Joint 2	Joint 3
LuGre	12.6312	10.1659	5.7290
Stribeck	12.6489	10.1978	6.3927
RBFNN	15.0079	12.2094	7.0691
RBFNN-H	10.6681	17.0561	5.2412
ParFam-V	12.8792	11.1926	6.0015
ParFam-H	8.1839	7.3823	4.1086

The reversal-window results highlight a sharper distinction between the models. ParFam-H obtains the lowest RMSE on all three joints and reduces reversal-region RMSE by 45.5%, 39.5%, and 41.9% compared with the standard RBFNN on Joints 1–3, respectively. This confirms that a purely velocity-based RBFNN tends to average over the multi-valued friction loop and smooth out the sharp torque transients needed near reversals. Meanwhile, RBFNN-H does not consistently improve over RBFNN, especially on Joint 2, demonstrating that hysteresis features alone are insufficient unless the model structure constrains them into a physically meaningful, compact relationship. In contrast, ParFam-H combines the memory coordinate with an explicit symbolic form, allowing it to reconstruct the reversal-induced torque peaks without

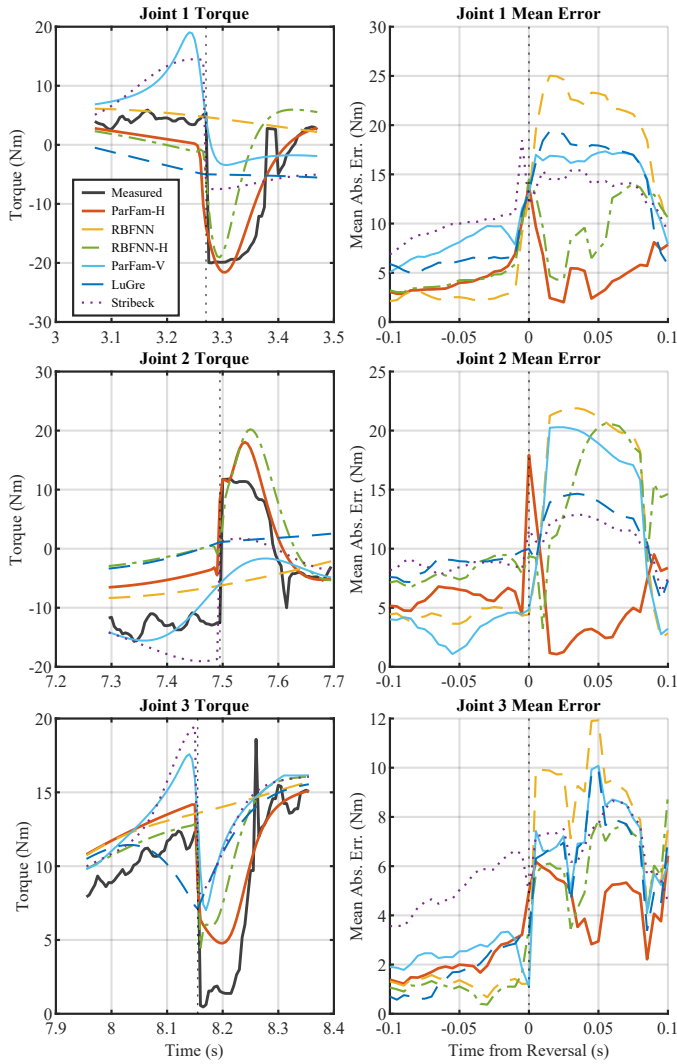


Fig. 5. Velocity-reversal residual-torque error analysis. The left panels show representative reversal segments, and the right panels report mean absolute error after aligning all detected velocity reversals. The standard velocity-only RBFNN tends to smooth through reversal-induced torque peaks, whereas ParFam-H better preserves the transient peak structure on average.

excessive oscillation.

F. Cross-Frequency Residual-Torque Generalization

As illustrated in Fig. 6, to assess out-of-distribution (OOD) residual-torque extrapolation, a variable-frequency Chirp signal (0.1–1.0 Hz) was applied, intentionally exceeding the 0.25 Hz training frequency. The data are segmented at 0.6 Hz ($t \approx 2.6$ s) into an in-distribution (low-frequency) region and an OOD (high-frequency) extrapolation region. The same Chirp dataset was evaluated for LuGre, Stribeck, RBFNN, RBFNN-H, ParFam-V, and ParFam-H. To keep Fig. 6 readable around the velocity-reversal segments, the plot highlights the hysteresis-augmented neural baseline (RBFNN-H) and the proposed method; the traditional physical models remained smoother but less accurate, consistent with the offline results in Table VIII.

In the low-frequency region, RBFNN-H and ParFam-H exhibit comparable precision (e.g., Joint 1 RMSE: 5.777

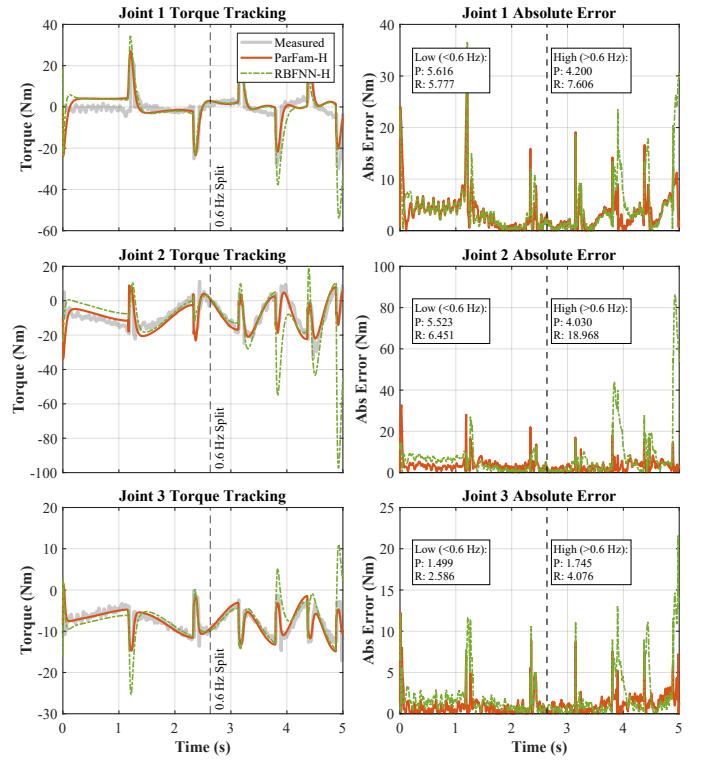


Fig. 6. Segmented residual-torque prediction analysis under a variable-frequency Chirp signal. The vertical dashed line at 0.6 Hz separates the in-distribution and out-of-distribution regions. ParFam-H maintains stability in the high-frequency region, whereas RBFNN-H degrades significantly.

N·m vs. 5.616 N·m), confirming the neural network’s strong in-distribution approximation ability. However, in the high-frequency region, their performances diverge. The prediction error of RBFNN-H increases significantly; its RMSE on Joint 2 increases from 6.451 N·m to 18.968 N·m (a nearly $3\times$ increase). This highlights a practical limitation of the tested black-box baseline: degraded extrapolation when encountering dynamic frequencies outside the training condition.

Conversely, ParFam-H maintains robustness in the high-frequency region, with its Joint 2 RMSE even decreasing from 5.523 N·m to 4.030 N·m. Observing the zero-velocity crossings, RBFNN-H exhibits non-physical overshoots and numerical oscillations, whereas the ParFam-H curve remains smooth and tightly tracks the measured values. This suggests that the evolved symbolic equations capture residual-torque structure that transfers better across the tested frequency range, enabling more stable predictions beyond the training operating point. Although the 5 s Chirp segment is sufficient to cover the full 0.1–1.0 Hz frequency sweep and multiple reversal events, it should be regarded as a frequency-extrapolation stress test rather than an endurance test. Longer-duration thermal and payload robustness tests are left for future work.

G. Closed-Loop Application Validation

To evaluate whether the identified residual-torque equation is useful beyond offline prediction, dynamic trajectory tracking experiments were conducted. Since the commercial UFACTORY-850 cobot, like most of its counterparts, restricts

direct joint torque control for safety, we developed a torque-error-based adaptive velocity feedforward controller.

The controller is designed around two practical constraints of commercial collaborative robots. First, the low-level torque-command interface is not available to users, so the compensation must be injected through the velocity-command interface. Second, the torque feedback estimated from motor current contains high-frequency noise; directly amplifying the torque error with fixed gains can therefore introduce velocity-command jitter. The core architecture utilizes the identified rigid-body and friction models to compute the theoretical total torque τ_d for the desired trajectory. The torque error $e_\tau = \tau_d - \tau_{meas}$ (where τ_{meas} is the noisy actual torque estimated from motor currents) is fed into a PD controller to generate a velocity compensation term:

$$\dot{q}_{comp} = K_p(|\dot{q}_d|)e_\tau + K_d(|\dot{q}_d|)\dot{e}_\tau. \quad (19)$$

To mitigate the amplification of high-frequency noise and maintain closed-loop stability, an adaptive gain scheduling mechanism is introduced (as depicted in the control block diagram, see Fig. 7). The PD gains decay exponentially with respect to the desired velocity $|\dot{q}_d|$:

$$K_{\{p,d\}}(|\dot{q}_d|) = K_{\{p,d\}}^{\min} + \left(K_{\{p,d\}}^{\max} - K_{\{p,d\}}^{\min}\right) e^{-\lambda|\dot{q}_d|}. \quad (20)$$

This scheduling is governed by dynamics dominance: near zero-velocity crossings, the system transitions through “stick-slip” phases where static friction and hysteresis peaks dominate the residual torque. Here, high gains are useful for compensating the reversal-induced torque mismatch and preventing severe phase lag. Conversely, at high speeds, motion is dominated by rigid-body inertia and smooth viscous friction. Maintaining high gains in this region would unnecessarily amplify high-frequency current sensor noise and potentially excite unmodeled flexible structural resonances. The final commanded velocity sent to the robot is $\dot{q}_{cmd} = \dot{q}_d + \dot{q}_{comp}$. Because direct torque commands are unavailable on the commercial platform, joint tracking error is used as the external application metric. A better residual-torque model reduces the torque mismatch that the inner velocity loop must absorb, which appears as smaller phase lag and lower absolute trajectory error. The nonzero errors in Fig. 8 therefore do not imply that the inner controller lacks basic friction compensation; rather, they quantify the residual dynamic mismatch under high-acceleration motion, current-estimation noise, and unmodeled flexibility.

Four compensation strategies were compared: Baseline (rigid-body only, no friction compensation), LuGre, RBFNN-H, and ParFam-H. The controller architecture, scheduling rule, and gain settings were kept fixed across the compensated cases; only the residual-torque feedforward model was replaced. As illustrated in Fig. 8, the uncompensated baseline suffers from severe phase lags and amplitude deviations during high-acceleration segments (e.g., Joint 2 at 1–3 s), confirming friction as a dominant dynamic disturbance. While all friction models improve tracking, ParFam-H achieves the tightest trajectory adherence. Notably, at $t \approx 8$ s on Joint 2, the absolute tracking error of the baseline and LuGre models exceeds 0.1 rad, whereas ParFam-H successfully suppresses the peak error to approximately 0.04 rad.

Quantitatively (Table X), ParFam-H yields the lowest RMSE across all three joints. On Joint 2, which experiences the most significant dynamic coupling and gravitational load, ParFam-H achieves an RMSE of 0.0377 rad. This represents a substantial trajectory-tracking RMSE reduction of 49.3% compared to the baseline (0.0744 rad), 38.4% against LuGre (0.0612 rad), and 21.0% against RBFNN-H (0.0477 rad). The remaining tracking error is mainly attributed to flexible structural resonances, current-to-torque estimation noise, temperature-dependent friction drift, and the limited bandwidth of the inner commercial controller. These factors are not explicitly modeled in the present feedforward formulation. The 10 s closed-loop trial contains repeated acceleration, deceleration, and reversal segments and is intended as a controlled real-time comparison among compensation methods; longer-duration robustness tests under changing thermal and payload conditions will be required before claiming endurance-level deployment. These results indicate that integrating the explicit ParFam-H model into real-time feedforward compensation improves dynamic trajectory tracking on the tested platform.

TABLE X
TRAJECTORY TRACKING RMSE UNDER DIFFERENT FRICTION
COMPENSATION STRATEGIES

Model	Joint 1	Joint 2	Joint 3
Baseline	0.0417	0.0744	0.0324
LuGre	0.0202	0.0612	0.0336
RBFNN-H	0.0204	0.0477	0.0267
ParFam-H	0.0177	0.0377	0.0233

To avoid relying only on a single full-trajectory RMSE value, the same 10 s closed-loop records were further divided into ten 1 s segments. In addition, tracking errors were evaluated inside ± 0.1 s windows around desired-velocity zero crossings. The file-to-model mapping follows the plotting script used to generate Fig. 8. As summarized in Table XI, ParFam-H achieves the lowest segment-averaged RMSE over all joints and also the lowest reversal-event RMSE for each joint. This confirms that its closed-loop advantage is maintained over multiple high-dynamic intervals and is not caused by a single favorable portion of the trajectory.

TABLE XI
SEGMENT-WISE AND REVERSAL-EVENT CLOSED-LOOP STATISTICS

Model	Seg. Avg.	Rev. J1	Rev. J2	Rev. J3
Baseline	0.0430	0.0341	0.0934	0.0254
LuGre	0.0335	0.0171	0.0748	0.0238
RBFNN-H	0.0280	0.0186	0.0614	0.0208
ParFam-H	0.0232	0.0161	0.0496	0.0194

VI. CONCLUSIONS AND FUTURE WORK

This paper presents HESR, a measurement-driven hysteresis-aware symbolic identification framework for residual torque in harmonic-drive robot joints. The proposed method first extracts current-based residual torque through rigid-body dynamic decoupling, then introduces a bounded hysteresis state to represent path-dependent memory near velocity reversals, and finally identifies compact analytical

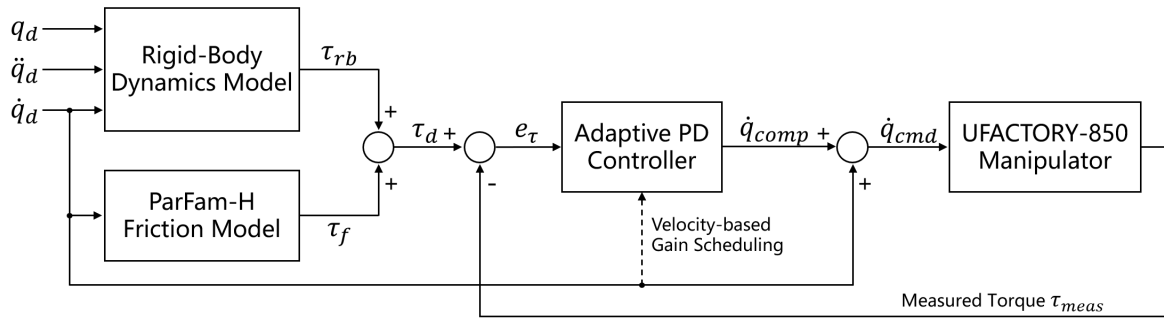


Fig. 7. Control block diagram of the torque-error-based adaptive velocity feedforward controller. The compensated residual-torque model and rigid-body dynamics generate the reference torque τ_d . The adaptive PD controller, scheduled by the reference velocity, compensates for residual errors.

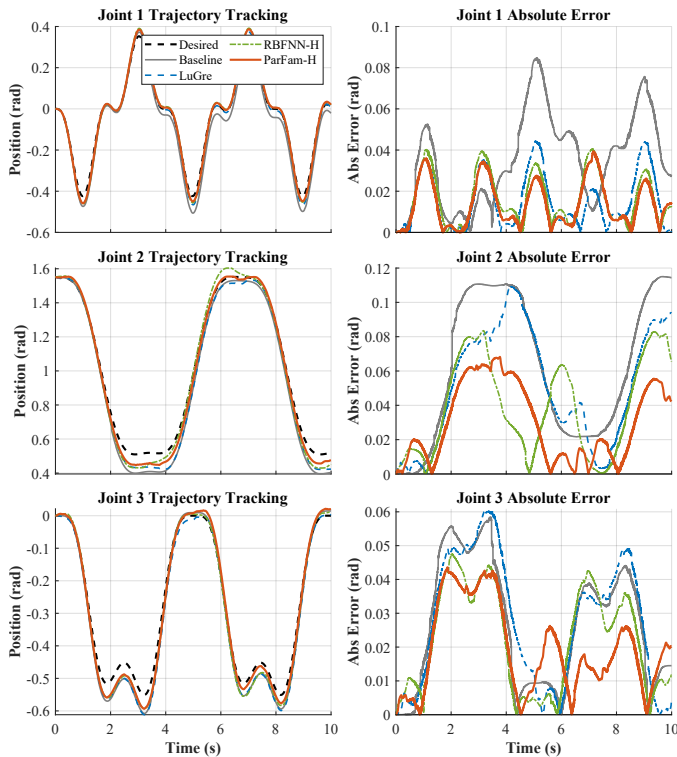


Fig. 8. Closed-loop application validation using dynamic trajectory tracking. ParFam-H reduces phase lag and peak tracking error when the identified residual-torque equation is used for feedforward compensation.

residual-torque equations using ParFam symbolic regression. In this way, the framework addresses both sides of the identification problem: obtaining reliable residual-torque measurements and converting them into interpretable equations.

Experimental results on a UFACTORY-850 manipulator show that the derived ParFam-H equations capture dynamic memory effects and provide the lowest offline residual-torque prediction RMSE among LuGre, Stribeck, standard RBFNN, RBFNN-H, and velocity-only symbolic baselines. In velocity-reversal windows, where hysteresis transients are most prominent, ParFam-H reduces RMSE by 45.5%, 39.5%, and 41.9% compared with the standard RBFNN on Joints 1–3, respectively. Threshold, filtering, and residual-level measurement-uncertainty sensitivity analyses further confirm that the mea-

surement pipeline is not dominated by a particular screening, smoothing, or perturbation setting. Variable-frequency Chirp tests demonstrate cross-frequency residual-torque generalization, avoiding the severe extrapolation degradation observed in the hysteresis-augmented neural baseline. Furthermore, integrating ParFam-H into a real-time feedforward controller reduced closed-loop trajectory-tracking RMSE by 49.3%, while segment-wise and reversal-event statistics confirm consistent improvements over high-dynamic intervals.

Beyond robotic friction, using a physics-informed hysteresis state variable provides a structured approach for symbolic regression to model systems with non-local memory effects by translating path-dependent dynamics into augmented Markovian features. The present study is limited to the first three joints of one UFACTORY-850 harmonic-drive platform, and the validation focuses on processed current-based torque measurements rather than direct torque-sensor measurements. Although the offline, reversal-window, Chirp, sensitivity, computation-time, and closed-loop results consistently support the proposed framework on this platform, broader generality under different robot structures, payloads, temperatures, and long-duration operating conditions remains to be verified. Thermal drift is particularly relevant for industrial robot accuracy because thermal-effect modeling has been shown to improve dynamic error compensation in robot measurement applications [32]. Future work will therefore extend the measurement campaign to additional harmonic-drive manipulators and operating conditions, investigate robustness against temperature drift and payload variation, and evaluate the method on other robotic systems governed by complex coupled dynamics, including hydraulic actuators and flexible-joint robots.

DATA AVAILABILITY

The processed experimental data, supplementary MATLAB scripts, and numerical result tables supporting the sensitivity and closed-loop analyses are provided in the accompanying reproducibility package. Robot network configuration and executable hardware-control scripts are excluded for safety and privacy reasons.

REFERENCES

- [1] J. J. Uicker, "On the dynamic analysis of spatial linkages using 4x4 matrices," Ph.D. dissertation, Northwestern Univ., Evanston, IL, USA, 1965.
- [2] J.-J. E. Slotine, "Putting physics in control: The example of robotics," *IEEE Control Syst. Mag.*, vol. 8, no. 6, pp. 12–18, Dec. 1988.
- [3] T. D. Tuttle and W. P. Seering, "A nonlinear model of a harmonic drive gear transmission," *IEEE Trans. Robot. Autom.*, vol. 12, no. 3, pp. 368–374, Jun. 1996.
- [4] H. Zhang, S. Ahmad, and G. Liu, "Modeling of torsional compliance and hysteresis behaviors in harmonic drives," *IEEE/ASME Trans. Mechatronics*, vol. 20, no. 1, pp. 178–185, 2015.
- [5] R. Dhaouadi, F. H. Ghorbel, and P. S. Gandhi, "A new dynamic model of hysteresis in harmonic drives," *IEEE Trans. Ind. Electron.*, vol. 50, no. 6, pp. 1165–1171, 2003.
- [6] J. Dong, J. Xu, Q. Zhou, J. Zhu, and L. Yu, "Dynamic identification of industrial robot based on nonlinear friction model and LS-SOS algorithm," *IEEE Trans. Instrum. Meas.*, vol. 70, pp. 1–12, 2021.
- [7] Y. Zhou, Z. Li, X. Zhang, *et al.*, "A semilinearized approach for dynamic identification of manipulator based on nonlinear friction model," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–20, 2024.
- [8] S. Wang, J. Wu, B. Li, *et al.*, "Adaptive acquisition and identification method of dynamic parameters based on least angle regression and its application in industrial robots," *IEEE Trans. Instrum. Meas.*, vol. 74, pp. 1–14, 2025.
- [9] T. Xun, J. Yang, and H. Ding, "Improving robotic grinding force control precision: Nonlinear friction compensation based on a novel continuous dynamic model," *IEEE/ASME Trans. Mechatronics*, 2024.
- [10] C. Canudas de Wit, H. Olsson, K. J. Åström, and P. Lischinsky, "A new model for control of systems with friction," *IEEE Trans. Autom. Control*, vol. 40, no. 3, pp. 419–425, Mar. 1995.
- [11] J. Swevers, F. Al-Bender, C. Ganseman, and T. Prajogo, "An integrated friction model structure with improved presliding behavior for accurate friction compensation," *IEEE Trans. Autom. Control*, vol. 45, no. 4, pp. 675–686, Apr. 2000.
- [12] M. R. Kermani, R. V. Patel, and M. Moallem, "Friction identification and compensation in robotic manipulators," *IEEE Trans. Instrum. Meas.*, vol. 56, no. 6, pp. 2346–2353, Dec. 2007.
- [13] M. K. Ciliz and M. Tomizuka, "Friction modelling and compensation for motion control using hybrid neural network models," *Eng. Appl. Artif. Intell.*, vol. 20, no. 7, pp. 898–911, 2007.
- [14] N. Hirose and R. Tajima, "Modeling of rolling friction by recurrent neural network using LSTM," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2017, pp. 6471–6478.
- [15] M. Lutter, C. Ritter, and J. Peters, "Deep Lagrangian networks: Using physics as model prior for deep learning," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2019.
- [16] M. Weiss, A. Pawluchin, J. H. Ewering, *et al.*, "Lagrangian neural network-based control: Improving robotic trajectory tracking via linearized feedback," *IEEE Robot. Autom. Lett.*, vol. 11, no. 3, pp. 2546–2553, 2026.
- [17] S.-M. Udrescu and M. Tegmark, "AI Feynman: A physics-inspired method for symbolic regression," *Sci. Adv.*, vol. 6, no. 16, p. eaay2631, 2020.
- [18] B. K. Petersen *et al.*, "Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021.
- [19] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proc. Natl. Acad. Sci.*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [20] P. Scholl *et al.*, "Interpretable robotic friction learning via symbolic regression," *arXiv preprint arXiv:2505.13186*, 2025.
- [21] N. Makke and S. Chawla, "Interpretable scientific discovery with symbolic regression: A review," *Artif. Intell. Rev.*, vol. 57, no. 1, 2024.
- [22] P. Scholl *et al.*, "ParFam: (Neural guided) symbolic regression via continuous global optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2025.
- [23] B. Armstrong-Helouvry, *Control of Machines with Friction*. Boston, MA, USA: Springer, 1991.
- [24] K. Guo, Y. Pan, and H. Yu, "Composite learning robot control with friction compensation: A neural network-based approach," *IEEE Trans. Ind. Electron.*, vol. 66, no. 10, pp. 7841–7851, Oct. 2019.
- [25] H. Yeo, J. Hong, T. Kong, *et al.*, "Mysteric-Net: MIMO hysteretic friction-aware Lagrangian-based network for legged robot," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2025, pp. 1932–1937.
- [26] J. Hwangbo, J. Lee, A. Dosovitskiy, *et al.*, "Learning agile and dynamic motor skills for legged robots," *Sci. Robot.*, vol. 4, no. 26, p. eaau5872, 2019.
- [27] H. Hu, Z. Shen, and C. Zhuang, "A PINN-based friction-inclusive dynamics modeling method for industrial robots," *IEEE Trans. Ind. Electron.*, vol. 72, no. 5, pp. 5136–5144, May 2025.
- [28] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019.
- [29] Z. Zhang and Z. Chen, "Modeling and control of robotic manipulators based on symbolic regression," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 5, pp. 2440–2450, 2023.
- [30] E. Lee, S. A. Moore, and B. Chen, "Sym2Real: Symbolic dynamics with residual learning for data-efficient adaptive control," *arXiv preprint arXiv:2509.15412*, 2025.
- [31] S. B. Šegota, V. Mrzljak, J. Prpić-Oršić, *et al.*, "Determining normalized friction torque of an industrial robotic manipulator using the symbolic regression method," *Industry 4.0*, vol. 8, no. 1, pp. 21–24, 2023.
- [32] R. Li and Y. Zhao, "Dynamic error compensation for industrial robot based on thermal effect model," *Measurement*, vol. 88, pp. 113–120, 2016.
- [33] M. Zhifu, L. Yong, D. Kenan, *et al.*, "A stepwise parameter identification strategy for probe-based robot on-machine measurement system calibration based on sensitivity analysis," *Measurement*, vol. 242, p. 115887, 2025.