Smart Materials and Structures

PAPER

Physics-guided modeling and semi-active control of magnetorheological actuators for quadruped locomotion

To cite this article: Erwin Ang Tien Yew *et al* 2026 *Smart Mater. Struct.* **35** 095008

Smart Materials and Structures

Check for updates

RECEIVED
25 June 2026

REVISED
22 August 2026

ACCEPTED FOR PUBLICATION
25 August 2026

PUBLISHED
8 September 2026

PAPER

Physics-guided modeling and semi-active control of magnetorheological actuators for quadruped locomotion

Erwin Ang Tien Yew[1,2,4], **Wenbo Li**[1,2,4], **Zhehan Tao**[1,2], **Zirui Chen**[1,2], **Zhongyuan Kong**[1,2], **Liang Zhang**[3,*] **and Shuaishuai Sun**[1,2,*]

[1] CAS Key Laboratory of Mechanical Behavior and Design of Materials, School of Engineering Sciences, University of Science and Technology of China, Hefei, Anhui 230026, People's Republic of China
[2] Institute of Humanoid Robots, School of Engineering Sciences, University of Science and Technology of China, Hefei, Anhui 230026, People's Republic of China
[3] School of Electrical Engineering and Automation, Anhui University, Hefei, Anhui 230601, People's Republic of China
[4] Authors contribute equally.
[*] Authors to whom any correspondence should be addressed.

E-mail: liangzhang@ahu.edu.cn and sssun@ustc.edu.cn

Keywords: magnetorheological robotics joint, magnetorheological actuator, actuator-aware reinforcement learning

Supplementary material for this article is available online

Abstract

Magnetorheological actuators (MRAs) provide current-dependent resistance that can be used for load support and damping in legged robots, but their nonlinear, hysteretic, and history-dependent torque response makes them difficult to model and schedule during locomotion. This paper presents MRA-aware semi-active locomotion (MRA-SAL), a physics-guided modeling and control framework for MR_Go-SAL, a quadruped robot equipped with calf-joint MRAs. MRA-SAL separates the semi-active function of the MRA into current-regulated stance holding and velocity-opposing disturbance damping. A current-aware Bouc–Wen model with a PG-RLSTM residual term predicts the dynamic MRA torque, while a pre-yield current–torque fit estimates the holding capacity. The resulting model is embedded into the Isaac Gym low-level torque loop for actuator-aware policy training, while the MRA-SAL scheduler is deployed on hardware. The PG-RLSTM achieves the lowest prediction error, with an average RMSE of 0.355 Nm on ten unseen actuator trajectories. Current-sweep disturbance tests show consistent attenuation trends in simulation and hardware: as current increases from 0 to 1.5 A, the minimum CoM height increases by 14.7% and 14.3%, respectively, while the peak downward velocity magnitude decreases by 19.3% and 35.4%. Locomotion experiments show COT reductions of 20.2–25.4% at 0.3 m s^{-1} and 21.3–24.9% at 0.5 m s^{-1} relative to baseline MR_Go. Across three repeated 300 m field tasks, MR_Go-SAL reduces mean battery state-of-charge depletion from 59.3$\pm$1.5 to 52.5 $\pm$ 2.0 percentage points relative to the matched MR_Go baseline. In the field-task thermal comparison, the mean temperature rise across the four calf motors decreases from 22.25 °C for MR_Go to 11.50 °C for MR_Go-SAL.

1. Introduction

Magnetorheological (MR) actuation [1–4] offers a promising route toward energy-efficient and impact-tolerant legged robots because the apparent resistance of the actuator can be regulated by magnetic field strength. Unlike conventional electric motors, an MR actuator does not behave as a simple commanded torque source. Its output depends on current, joint displacement, joint velocity, recent motion history, and whether the MR fluid is operating in a pre-yield or post-yield regime. These properties allow the actuator to provide load-bearing resistance during slow stance and damping during impact or disturbance, but they also make the actuator difficult to represent and schedule during continuous locomotion.

For quadruped locomotion, this mode-dependent behavior is especially important at the calf joints, where stance loading and impact-related torque variation are large. During loaded near-zero-velocity stance, the useful role of the magnetorheological actuator (MRA) is to share the holding load with the electric motor. During disturbance-induced joint motion, the useful role is instead to oppose harmful motion and dissipate energy. Therefore, the control problem is not only to predict MRA torque, but also to decide when the same semi-active device should support, damp, or release the leg.

This problem is related to a broader theme in legged robotics: when the actuator contains non-rigid or material-dependent dynamics, its physical behavior must be represented explicitly. Model-based locomotion provided one route for exploiting such actuator behavior, especially for legged robots with elastic elements, because compliant dynamics can be described by structured mechanical models and used for planning and control. Representative examples include the SLIP-inspired study by Lakatos *et al* [5], which embodied SLIP-like dynamics in a compliantly actuated quadruped to realize dynamic locomotion gaits, as well as the PEA-based jumping frameworks of Ding *et al* [6, 7], where actuated-SLIP or SLIP-inspired reduced-order models were combined with trajectory optimization and landing regulation to generate robust jumping, landing, and recovery motions.

While these model-based studies demonstrate the value of exploiting compliant actuator dynamics through structured mechanical models, recent learning-based locomotion studies have shown that actuator-specific dynamics can also be incorporated into data-driven training and control. Hwangbo *et al* [8] introduced actuator-network modeling for sim-to-real transfer, Raffin *et al* [9] used model-free reinforcement learning to improve locomotion on an SEA quadruped, while Bjelonic *et al* [10] co-optimized parallel-elastic knee actuation and locomotion control on ANYmal. These studies demonstrate that non-rigid actuation can be beneficial when its mechanical behavior is properly represented in the controller or simulation environment.

Actuator-aware simulation is one practical way to represent such behavior during policy training and evaluation. Lee *et al* [11] showed that incorporating actuator models into the reinforcement-learning environment can improve transfer for hydraulic quadrupeds, and Yuryev [12] demonstrated that learned actuator-force models can help represent difficult-to-model tendon-driven systems. These works motivate the use of measured actuator data and temporal models when the actuator response cannot be captured accurately by a simple analytical torque equation.

To model the nonlinear and history-dependent behavior of MRAs, both physics-based and data-driven approaches have been widely explored. Physics-based models such as Bouc–Wen provide structured hysteresis representation [13], and various modified Bouc–Wen formulations have been proposed to further improve modeling accuracy for different magnetorheological dampers [14, 15]. However, these models are still limited in capturing current-dependent and trajectory-specific variations. In contrast, purely data-driven neural networks can approximate complex nonlinear mappings but often lack physical consistency and generalization under unseen operating conditions. To address these limitations, recent studies have increasingly adopted hybrid modeling strategies that integrate structured hysteresis models with physics-informed learning frameworks such as physics-informed neural networks [16, 17] and physics-informed residual learning networks [18], building upon classical hysteresis formulations such as Bouc–Wen models [19]. Complementary studies on data-driven MR modeling further include LSTM-based force prediction approaches [20], improved neural architectures for MR brake systems [21], and optimization-driven neural network enhancement methods [22, 23]. However, most of these approaches remain focused on isolated damper-level modeling and do not consider the integration of magnetorheological actuation into closed-loop robotic locomotion systems, where actuation scheduling and contact-dependent behavior are critical.

In this work, we present *MRA-aware semi-active locomotion* (MRA-SAL), a physics-guided semi-active locomotion framework for the MR_Go quadruped. Figure 1 summarizes the overall pipeline, from data-driven MRA torque modeling and actuator-aware policy training to locomotion-state-dependent semi-active scheduling. Unlike our previous MR_Go study [24], in which MR actuation was not used during continuous locomotion, the present work connects experimentally identified actuator behavior with actuator-aware policy training and locomotion-state-dependent semi-active control. The principal novelty lies in integrating these elements within a unified locomotion framework rather than in any individual component in isolation.

The main contributions of this paper are as follows:

- We formulate MRA-SAL, a semi-active control framework that uses calf-joint MRAs for both quasi-static stance support and disturbance-triggered velocity-opposing damping.

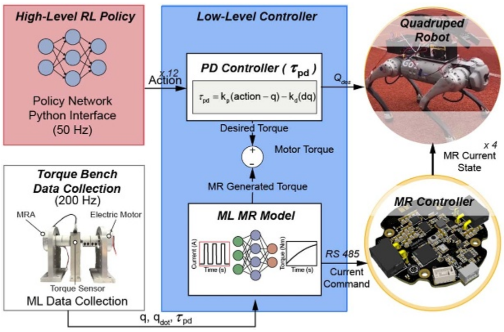

Figure 1. Overall MRA-SAL pipeline. Real MRA torque data are used to train a Bouc–Wen-informed PG-RLSTM, which is embedded into the reinforcement-learning simulation loop. The resulting actuator-aware environment is used to evaluate a semi-active scheduler that separates quasi-static stance support from disturbance-triggered velocity-opposing damping on MR_Go-SAL.

- We develop a current-aware Bouc–Wen-informed PG-RLSTM torque model that combines a physics-based hysteresis baseline with a recurrent residual term to represent current- and history-dependent MRA behavior.
- We integrate the learned MRA model into an actuator-aware locomotion simulation and validate the complete framework through actuator-level prediction tests, simulation current-sweep studies, and hardware experiments measuring calf torque relief, disturbance response, cost of transport (COT), electrical power consumption, and actuator temperature.

2. MR_Go platform and MRA operating characteristics

2.1. MR_Go-SAL quadruped platform

MR_Go-SAL is developed from our previous MR-actuated quadruped platform, in which magnetorheological actuation was integrated into the leg joints to improve load support, impact tolerance, and energy efficiency [24]. In the present study, the four calf-joint MRAs are used as the main semi-active elements because the calf joints experience large stance loads and strong impact-related torque variation during locomotion. Compared with the previous hardware configuration, MR_Go-SAL is equipped with four independent current controllers, enabling joint-level current modulation and feedback regulation of each calf-joint MRA during walking and disturbance response.

2.2. MRA operating principle

The MRA generates semi-active resistance through the field-dependent rheology of the enclosed magnetorheological fluid [25, 26]. When current is applied to the coil, magnetic particles form chain-like structures in the fluid, increasing the shear resistance and the torque required to rotate the actuator. Because this resistance originates from the material state rather than from active motor power, the MRA is suitable for load sharing and energy dissipation, but it cannot be treated as a conventional commanded torque source. Two operating regimes are important for locomotion. In the pre-yield regime, the MR particle-chain structure can resist small displacement under load, so the actuator behaves as a current-dependent holding or braking element. This regime is useful during loaded, low-velocity stance, where the calf motor would otherwise need to continuously produce proportional-control torque to maintain posture. In the post-yield regime, the applied load exceeds the current-dependent yield capacity and the actuator response becomes motion-dependent. This regime is useful when a disturbance or impact causes the loaded calf joint to move away from its target, because the MRA can provide velocity-opposing resistance and attenuate the motion. Detailed actuator design and characterization are reported in [27]. In this paper, we use the operating facts most relevant to locomotion: the MRA provides current-dependent pre-yield braking torque up to approximately 26.15 Nm at 1.5 A, and the current–torque relation follows a saturating trend. These properties motivate a locomotion controller that can use the MRA as a holding element during quasi-static stance and as a damping element during disturbance-induced motion.

2.3. Problem formulation

Our previous MR_Go work [24] demonstrated the hardware benefit of MRA integration through improved payload support, impact tolerance, energy efficiency, and thermal stability. However, the MR actuation in that study was mainly activated through task-specific finite-state-machine states, such as landing-impact mitigation and stationary load support. During continuous locomotion, the robot still relied primarily on a conventional locomotion controller, and the MRA behavior was not explicitly represented inside the learning-based simulation environment. Therefore, the previous work verified the usefulness of the hardware, but did not address how the current- and state-dependent MRA torque should be modeled and scheduled during repeated gait cycles.

In the present work, the calf-joint torque is treated as the combined effect of the electric motor and the semi-active MRA [28]. For the ith calf joint, the total joint torque can be written as

$$\tau_i = \tau_i^{\mathrm{PD}} + \tau_{\mathrm{MR},i}, \tag{1}$$

where τ_i^{PD} is the motor-side torque generated by the low-level PD controller and $\tau_{\mathrm{MR},i}$ is the additional semi-active torque generated by the MRA. For the calf-joint implementation used in this study, the PD torque is

$$\tau_i^{\mathrm{PD}} = K_{p,i}\left(q_i^* - q_i\right) - K_{d,i}\dot{q}_i\,, \tag{2}$$

where q_i^* is the policy-provided joint target, q_i and $\dot{q}_i$ are the measured joint position and velocity, and $K_{p,i}$ and $K_{d,i}$ are the motor-side feedback gains.

The central modeling problem is therefore the representation of $\tau_{\mathrm{MR},i}$. Unlike τ_i^{PD}, the MRA torque cannot be obtained directly from a commanded torque input. It depends on coil current, joint displacement, joint velocity, recent motion history, and whether the actuator is operating in a pre-yield holding regime or a post-yield motion-dependent regime. Thus, a fixed damping coefficient or a purely static current–torque map is insufficient for representing the MRA during locomotion.

This leads to two coupled questions addressed in the following sections. First, how should the semi-active MRA contribution be enabled or released during locomotion so that it supports loaded stance and damps harmful disturbance-induced motion without resisting useful leg recovery? Second, how can the dynamic MRA torque contribution be represented accurately enough for actuator-aware locomotion simulation? Section 3 addresses the first question using the MRA-SAL semi-active scheduler, while section 4 addresses the second question using a current-aware Bouc–Wen baseline with a residual LSTM.

3. MRA-SAL semi-active torque formulation

3.1. Semi-active decomposition of MRA torque

During locomotion, the same calf-joint MRA is used for two physically different functions. When a loaded calf joint is disturbed and moves away from its policy target, the MRA should behave as a velocity-opposing damper. When the calf joint is loaded, close to stationary, the MRA should behave as a quasi-static holding element that shares the stance load with the electric motor. Therefore, the MRA torque contribution is written as

$$\tau_{\mathrm{MR},i} = m_i^{\mathrm{dyn}}\tau_{\mathrm{MR},i}^{\mathrm{dyn}} + m_i^{\mathrm{hold}}\tau_{\mathrm{hold},i}, \tag{3}$$

where m_i^{dyn} and m_i^{hold} are binary gates for the dynamic damping and quasi-static holding branches, respectively. The term $\tau_{\mathrm{MR},i}^{\mathrm{dyn}}$ represents the motion-dependent damping torque available from the MRA, while $\tau_{\mathrm{hold},i}$ represents the pre-yield holding resistance available during loaded near-zero-velocity stance. The two branches are defined by the following local gating conditions.

3.2. Disturbance-triggered dynamic damping

The dynamic damping branch is enabled when a loaded calf joint moves in a direction that increases its tracking error. Let

$$e_i = q_i^* - q_i \tag{4}$$

denote the calf-joint tracking error, where q_i^* is the policy-provided joint target and q_i is the measured joint position. The dynamic damping gate is defined as

$$\begin{aligned} m_i^{\mathrm{dyn}} = {} & \mathbf{1}\left(\|\mathbf{f}_{\mathrm{foot},i}\| > f_{\mathrm{th}}\right)\mathbf{1}\left(|e_i| > \epsilon_{\mathrm{e}}\right) \\ & \times \mathbf{1}\left(|\dot{q}_i| > \epsilon_{\mathrm{v}}\right)\mathbf{1}\left(e_i\dot{q}_i < 0\right), \end{aligned} \tag{5}$$

where $\mathbf{1}(\cdot)$ is the indicator function. The first condition ensures that the leg is loaded, the second and third conditions avoid unnecessary activation for small tracking errors or small velocities, and the final condition identifies target-escaping motion. Since $e_i = q_i^* - q_i$, the condition $e_i\dot{q}_i < 0$ indicates that the measured joint motion is increasing the tracking error.

When the dynamic branch is active, the MRA torque is applied opposite to the measured joint velocity:

$$\tau_{\mathrm{MR},i}^{\mathrm{dyn}} = -\mathrm{sgn}\left(\dot{q}_i\right)\left|\hat{\tau}_{\mathrm{dyn},i}\right|, \tag{6}$$

where $\hat{\tau}_{\mathrm{dyn},i}$ is the predicted magnitude of the available dynamic MRA torque. Its prediction model is described in section 4. This sign assignment guarantees that the dynamic branch does not inject positive mechanical work:

$$\tau_{\mathrm{MR},i}^{\mathrm{dyn}}\dot{q}_i = -|\dot{q}_i|\left|\hat{\tau}_{\mathrm{dyn},i}\right| \leqslant 0. \tag{7}$$

Thus, the dynamic branch acts as a semi-active damping element during harmful target-escaping motion. In hardware, the same gate is used to enable MR current in the physical actuator, and the corresponding resistive torque is generated mechanically by the MRA.

The nominal dynamic-gate thresholds were $f_{\mathrm{th}} = 5$ N, $\epsilon_{\mathrm{e}} = 0.05$ rad, and $\epsilon_{\mathrm{v}} = 0.20$ rad s^{-1}. Initial values were selected from nominal locomotion simulation data by separating loaded stance from swing-phase sensor noise and excluding small tracking and velocity fluctuations. They were subsequently refined in hardware experiments and fixed before the reported comparisons.

A one-factor-at-a-time analysis under matched disturbances showed that, for contact thresholds of 3, 5, 7, and 13 N, the reduction in CoM drop remained between 13.47% and 15.25%, while the motor-PD-torque reduction remained between 10.20% and 10.50%. Increasing ϵ_{e} from 0.01 to 0.09 rad reduced the CoM-drop improvement from 15.33% to 12.61% and the motor-torque reduction from 10.32% to 8.82%. Increasing ϵ_{v} from 0.10 to 0.30 rad s^{-1} reduced the CoM-drop improvement from 16.25% to 13.19%. Therefore, the dynamic gate was comparatively insensitive to the contact threshold over the tested range, whereas ϵ_{e} and ϵ_{v} more directly affected its activation. Larger error and velocity thresholds made the gate more selective, reducing the available damping interval and weakening disturbance attenuation. The nominal values therefore balance rejection of small tracking and velocity fluctuations against timely activation of MRA damping.

3.3. Quasi-static stance support

The quasi-static support branch is used during loaded stance when the calf joint is close to stationary. In this regime, the MRA is not used to dissipate dynamic motion. Instead, it provides pre-yield holding resistance so that the electric motor can maintain the stance posture with a reduced proportional gain.

The holding gate is defined as

$$\begin{aligned} m_i^{\mathrm{hold}} = {} & \mathbf{1}\left(\|\mathbf{f}_{\mathrm{foot},i}\| > f_{\mathrm{th}}\right)\mathbf{1}\left(|\dot{q}_i| < \epsilon_{\mathrm{hold}}\right) \\ & \times \mathbf{1}\left(m_i^{\mathrm{dyn}} = 0\right). \end{aligned} \tag{8}$$

The first two conditions identify loaded, near-stationary stance, while the final condition prevents simultaneous activation of the holding and dynamic damping branches.

When the holding branch is active, the motor-side proportional gain is reduced from the nominal value to a low-gain setting:

$$K_{p,i}^{\mathrm{eff}} = \begin{cases} K_{p,i}^{\mathrm{low}}, & m_i^{\mathrm{hold}} = 1, \\ K_{p,i}^{\mathrm{base}}, & m_i^{\mathrm{hold}} = 0. \end{cases} \tag{9}$$

The resulting calf-joint torque under support mode can be expressed as

$$\tau_i = K_{p,i}^{\mathrm{eff}}\left(q_i^* - q_i\right) - K_{d,i}\dot{q}_i + \tau_{\mathrm{MR},i}. \tag{10}$$

Therefore, during quasi-static support, part of the stance load is carried by the MRA rather than entirely by the electric motor. When the joint begins moving again or the contact condition is no longer satisfied, the holding gate is disabled, the MRA is released, and the nominal motor gain is restored so that the leg can retract and reposition freely.

Based on this operating logic, the nominal holding threshold was $\epsilon_{\mathrm{hold}} = 0.20\ \mathrm{rad\,s^{-1}}$. This value was initially selected from nominal locomotion simulation data to identify near-stationary stance, subsequently refined in hardware experiments, and fixed before the reported comparisons. Increasing ϵ_{hold} extends the holding interval, but an excessively large value can activate support during finite joint motion, where the reduced motor-side K_p may weaken joint regulation and compromise stability. Conversely, decreasing ϵ_{hold} shortens the holding interval, reduces the opportunity for MRA load sharing, and may increase COT. The selected value therefore balances stable joint regulation with COT reduction under the tested gait and operating conditions.

4. Physics-guided representation of MRA torque

This section describes how the available torque in each branch is represented. The dynamic branch requires a history-dependent torque model because the post-yield MRA response depends on joint motion, current, and recent actuator history. The holding branch is represented using the experimentally fitted pre-yield current–torque relation, because its function is to provide quasi-static resistance during loaded near-zero-velocity stance.

4.1. Measured MRA dataset and learning samples

Actuator data are collected on a single-joint torque test bench to characterize the current- and motion-dependent response of the calf-joint MRA. An external motor prescribes the joint motion, the MRA generates current-dependent resistive torque, and a torque sensor measures the actuator output. The MR current controller applies both constant and time-varying current profiles, while encoder position, encoder velocity, command current, and output torque are recorded at 200 Hz.

The recorded variables include timestamp, joint displacement, joint velocity, coil current, and measured output torque, forming the dataset

$$\mathcal{D} = \left\{\left(t, \theta(t), \omega(t), I(t), \tau(t)\right)\right\}_{t=1}^{N}. \tag{11}$$

Because the dynamic MRA torque depends on recent motion and current history, each recurrent-model input contains the most recent 50 samples of joint displacement, joint velocity, and coil current, corresponding to a 0.25 s temporal history at the 200 Hz sampling rate. During offline training, these temporal-history samples are generated from the recorded trajectories, and the recurrent-model input and target are defined as

$$\mathbf{X}_{\mathrm{MR}}(t) = [\theta, \omega, I]_{t-49:t}, \qquad y(t) = \tau(t). \tag{12}$$

For the MLP baseline, the joint displacement, joint velocity, and coil current from three consecutive samples are flattened into a nine-dimensional input vector. The network contains two fully connected hidden layers with 64 units and Tanh activation, followed by a scalar output layer that predicts the actuator torque. The LSTM baseline uses the same three input features over a 50-sample sequence. It consists of one unidirectional LSTM layer with 64 hidden units followed by a $64 \to 1$ fully connected layer that predicts the actuator torque from the final recurrent output.

To cover the actuator operating range relevant to quadruped locomotion, the dataset includes randomized current-magnetization step responses, multi-amplitude sinusoidal motion from 0.5 to 2.0 Hz, and representative calf-joint trajectories collected from robot motion. The model-fitting data cover 0–1.5 A in 0.5 A increments; the 1.2 A condition was not included. The experiments cover angular velocities up to about 10 rad $\mathrm{s^{-1}}$ and displacements up to about 1 rad.

4.2. Current-Aware Bouc–Wen baseline

The Bouc–Wen model is used as a compact physics-inspired baseline for the dynamic MRA torque because it represents hysteresis through an internal memory state. In this work, the Bouc–Wen component is not used only as a diagnostic comparison; it is embedded directly into the forward prediction of the proposed residual model. Its role is to capture the structured part of the current-dependent hysteretic response, while the recurrent residual branch learns the remaining modeling error.

A conventional first-order Bouc–Wen element represents hysteresis using an internal memory variable z and can be written as

$$\begin{aligned} \tau_{\mathrm{conv}} &= D + B\dot{q} + Cq + \kappa z, \\ \dot{z} &= A\dot{q} - \beta|\dot{q}|z - \gamma\dot{q}|z|, \end{aligned} \tag{13}$$

where κ is a constant hysteresis gain and A, β, and γ shape the hysteresis evolution. This form is useful for rate-dependent hysteresis, but it does not explicitly account for coil-current modulation of MR resistance. It also does not represent the speed-weakening trend observed in the measured MRA data, where the history-dependent torque contribution becomes attenuated at higher joint velocity.

To represent the MR_Go calf-joint MRA, the baseline is modified to include passive off-current drag, current-dependent resistance, and speed-weakening of the hysteretic contribution. The resulting current-aware Bouc–Wen torque is

$$\tau_{\mathrm{BW}} = D + B\dot{q} + Cq_c + I_1 I + \frac{K_0 + K_1 I}{1 + \alpha|\dot{q}|} z, \tag{14}$$

where q_c is the offset-corrected joint displacement, $\dot{q}$ is joint velocity, I is coil current, and z is the hysteresis memory state. The identified coefficients used in the current-aware Bouc–Wen baseline were $q_{\mathrm{offset}} = 2.39638101 \times 10^{-4}$, $q_{\mathrm{center}} = -1.51224291 \times 10^{-17}$, $D = -0.01917639$, $B = 0.1013626$, $C = 0.2160456$, $I_1 = -7.681742 \times 10^{-4}$, $K_0 = 0.3992754$, $K_1 = 10.413315$, $\alpha = 0.7777051$, $\beta = 16.50618$, and $\gamma = -10.53036$. The state evolves according to

$$\dot{z} = \dot{q} - \beta|\dot{q}|z - \gamma\dot{q}|z|. \tag{15}$$

Compared with equation (13), the modified baseline introduces three changes. First, the constant hysteresis gain κ is replaced by the current-dependent gain $K_0 + K_1 I$, allowing the same memory state to generate different torque levels under different coil currents. Second, the direct current term $I_1 I$ captures current-dependent torque bias not explained by hysteresis alone. Third, the denominator $1 + \alpha|\dot{q}|$ attenuates the hysteretic contribution at higher joint velocity. The coefficients are identified from the training data and kept fixed when training and deploying the PG-RLSTM. At $I = 0$, the model still predicts passive mechanical and fluid drag, which is retained in the deployed actuator model. The current-aware Bouc–Wen coefficients were obtained by numerically fitting the model to measured torque data collected at 0, 0.5, 1.0, and 1.5 A. Data collected at 1.2 A were excluded from parameter identification and used only for interpolation evaluation.

4.3. Bouc–Wen-informed PG-RLSTM for dynamic torque

Although the current-aware Bouc–Wen baseline captures the main hysteretic and current-dependent trend, the measured MRA response also contains effects that are difficult to represent with a compact analytical equation, including current-response delay, frictional asymmetry, sensor bias, and trajectory-dependent nonlinearities. Therefore, a Bouc–Wen-informed physics-guided residual LSTM (PG-RLSTM) is used to represent the dynamic MRA torque. The analytical branch provides the structured baseline response, while the recurrent branch learns the residual torque error from temporal actuator data.

For each input window $\mathbf{X}_{\mathrm{MR}}(t)$, the Bouc–Wen branch first produces the baseline prediction $\tau_{\mathrm{BW}}(t)$. The recurrent branch then predicts the remaining torque error, giving the final dynamic MRA torque prediction

$$\hat{\tau}_{\mathrm{dyn}}(t) = \tau_{\mathrm{BW}}(t) + \Delta\tau_{\text{PG-RLSTM}}(t), \tag{16}$$

where $\Delta\tau_{\text{PG-RLSTM}}(t)$ is trained using the residual target

$$\Delta\tau^{*}(t) = \tau_{\mathrm{measured}}(t) - \tau_{\mathrm{BW}}(t). \tag{17}$$

Physics guidance is introduced structurally through the current-aware Bouc–Wen baseline, whose coefficients remain fixed during residual training. No explicit physics-equation or passivity term is included in the training loss. The residual branch is optimized using the mean-squared error between its output and the residual target in equation (17). Because the learned residual is unconstrained in sign, it can locally increase or decrease the Bouc–Wen prediction. Consequently, dissipative behavior is not assumed from the predictor itself but is enforced separately through the semi-active torque projection in equations (6) and (7).

This residual formulation preserves the current-aware Bouc–Wen structure while allowing the LSTM to correct systematic errors associated with transient and history-dependent actuator behavior. In the MRA-SAL damping branch, the predicted dynamic torque magnitude $|\hat{\tau}_{\mathrm{dyn}}|$ is used, and the final torque direction is assigned opposite to joint velocity according to equation (6). This ensures that the learned dynamic torque model is used in a semi-active, velocity-opposing manner.

Table 1. Architecture and training configuration of theresidual LSTM.

Item	Configuration
Input features	Joint displacement q, joint velocity $\dot{q}$, and coil current I
Sequence length	50 samples
Network structure	One unidirectional LSTM layer followed by one fully connected output layer
Hidden units	64
Output layer	$64 \rightarrow 1$; output represents the residual torque$\Delta\tau_{\mathrm{LSTM}}$
Optimizer	Adam
Learning rate	1×10^{-3}
Dataset size	665,442 time samples
Dataset division	Trajectory-level: 80% training; 20% internal validation

Table 1 summarizes the residual-LSTM architecture and training settings. The Bouc–Wen internal hysteresis state was integrated with a 0.005 s time step and bounded to $[-5, 5]$ to prevent numerical instability. The MLP and LSTM baselines directly predict the measured actuator torque and are optimized using the mean-squared-error loss. They use the same measured dataset, trajectory-level training–validation division, training-set input standardization, Adam optimizer, and learning rate of 1×10^{-3} as the PG-RLSTM. The principal difference is the prediction target: the two baselines predict the total actuator torque, whereas the PG-RLSTM predicts the residual torque relative to the fixed current-aware Bouc–Wen model.

The input features were standardized using a StandardScaler fitted exclusively on the training set, $x_{\mathrm{norm}} = (x - \mu_{\mathrm{train}})/\sigma_{\mathrm{train}}$, and the same transformation was applied to the internal-validation, playback, and deployment data. The target torque remained in Nm. Before constructing the temporal inputs, the source actuator trajectories were randomly divided at the trajectory level into 80% training trajectories and 20% internal-validation trajectories. Overlapping 50-sample sliding windows were subsequently generated within each trajectory, and all windows derived from a given trajectory remained in the same subset. Thus, no trajectory, raw time sample, or overlapping window was shared between the training and internal-validation subsets. Generalization was evaluated separately on ten complete trajectories excluded from both model fitting and internal validation; the temporal windows for these trajectories were constructed independently for the final evaluation.

4.4. Pre-yield holding torque representation

The quasi-static holding branch uses the pre-yield resistance of the MRA. In this regime, the MR particle-chain structure can resist small joint displacement under load, allowing the MRA to share stance support with the electric motor. Unlike the dynamic branch, this function does not require a history-dependent recurrent model because the holding capacity is mainly determined by the applied current under near-zero-velocity conditions.

The available holding torque is represented using the experimentally measured pre-yield current–torque relation,

$$\tau_{\mathrm{hold}}(I) \approx \tau_y(I) = \frac{T_{\max} I^{\,n}}{K^n + I^{\,n}}, \tag{18}$$

where $T_{\max}$ is the saturation torque and K and n shape the current–torque curve. From the pre-yield braking-torque fit, $T_{\max} = 26.59697$ Nm, $K = 0.51272$, and $n = 2.87967$.

This relation provides the current-dependent capacity used by the quasi-static holding branch, whereas the PG-RLSTM model represents the motion-dependent torque used by the dynamic damping branch. Equation (18) defines the maximum available holding capacity. In the simulated joint dynamics, the holding torque compensates for the reduction in motor-side proportional gain and is bounded by this capacity:

$$\begin{aligned}\tau_{\mathrm{hold},i} = \mathrm{clip}\Big(&\left(K_{p,i}^{\mathrm{base}} - K_{p,i}^{\mathrm{low}}\right)\left(q_i^* - q_i\right), \\ &-\tau_y(I_i),\ \tau_y(I_i)\Big).\end{aligned} \tag{19}$$

This term represents a passive reaction torque rather than an independently commanded torque source. When the required compensation is within $\tau_y(I_i)$, the MRA provides only the required holding torque. When it exceeds this capacity, the MRA provides its maximum resistance in the required direction, and the remaining unsupported torque causes joint motion.

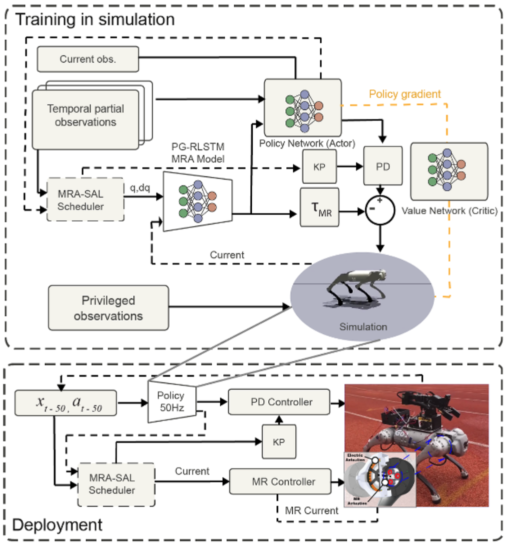

Figure 2. Training and deployment architecture of the proposed MRA-SAL framework.

Together, the PG-RLSTM dynamic torque model and the pre-yield holding relation provide the two torque representations required by MRA-SAL. The next section describes how these representations are embedded into the actuator-aware simulation and hardware deployment pipeline.

5. MR-aware simulation and hardware deployment

5.1. MR-aware simulation environment

Figure 2 summarizes how the learned MRA torque model is connected to simulation and hardware deployment. In simulation, the locomotion policy provides joint targets to the low-level PD controller. The motor-side PD torque is computed from the joint tracking error, while the MRA-SAL scheduler determines whether the dynamic damping or quasi-static holding branch should be enabled at each calf joint. When the dynamic branch is active, the PG-RLSTM model provides the available MRA torque magnitude from the recent displacement, velocity, and current history. When the holding branch is active, the pre-yield current–torque relation provides the available quasi-static holding resistance.

The locomotion policy is updated at 50 Hz, while low-level torque computation and physics stepping are performed at 200 Hz with a 5 ms step. The PG-RLSTM torque model is evaluated in the same low-level loop so that the simulated calf-joint dynamics include the measured current- and history-dependent MRA response. The gait frequency is set to 1.5 Hz for policy training and evaluation, matching the low-speed walking regime considered in this study. This setting provides sufficiently long stance intervals for the quasi-static MRA support branch to operate when the calf joint is loaded and near zero velocity.

To expose the policy and evaluation rollouts to different MR resistance levels, the parallel simulation environments are divided evenly across fixed current conditions,

$$I_e \in \{0,\ 0.5,\ 1.0,\ 1.5\}\ \text{A}, \qquad g(e) = \left\lfloor \frac{4e}{N_{\text{env}}} \right\rfloor, \tag{20}$$

where e is the environment index, N_{env} is the number of parallel environments, and $g(e)$ assigns each environment to one of the four current groups. The 0 A group represents the off-current MRA condition, including passive mechanical and fluid drag, while the nonzero-current groups represent increasing MR resistance. This implementation allows the locomotion policy and evaluation rollouts to account for both active motor control and semi-active MRA resistance.

5.2. Hardware deployment of MRA-SAL

In hardware deployment, the trained policy provides joint targets to the motor PD controller, while the MRA-SAL scheduler runs locally for each calf joint using joint state and contact feedback. The dynamic damping gate is used to enable MR current when the loaded calf joint moves away from its target. In this case, the MRA generates velocity-opposing resistance mechanically, and no additional software torque is injected into the motor command.

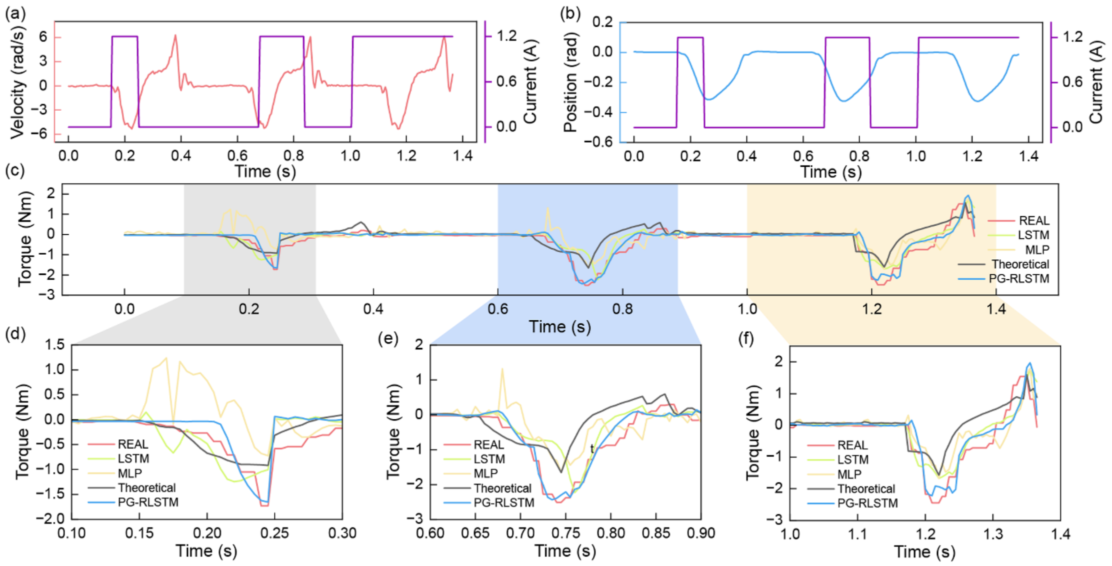

Figure 3. Representative torque-prediction results on an unseen MRA trajectory. Measured joint position, joint velocity, and current are fed to the MLP, LSTM, and PG-RLSTM models, and the predicted torque is compared with the measured MRA torque under three selected MR-current-on sections.

For quasi-static support, the holding gate enables the MRA when the calf joint is loaded and near stationary, while the dynamic damping branch is inactive. During this interval, the motor-side proportional gain is reduced to the low-gain setting, and the MRA provides pre-yield holding resistance to share the stance load. For hardware repeatability, the support current is not varied continuously during every stance interval. MRF yield stress and viscosity depend on temperature and operating history, the available holding capacity may vary during prolonged operation. Online temperature compensation is not included in the present implementation; therefore, a fixed support current is selected from the pre-yield current–torque relation in equation (18) using a safety factor of 1.2 to provide the required torque margin and accommodate possible variations in holding capacity. This setting maintained stable hardware operation during field tests lasting up to 25 min.

During hardware locomotion, the yield state is not measured directly. Instead, the controller infers loss of quasi-static holding from the joint proprioceptive response. When an applied load exceeds the available yield capacity, the resulting slip increases the joint velocity and tracking error. Once the near-zero-velocity condition is no longer satisfied, the holding gate is disabled and the nominal proportional gain is restored according to equation (9), allowing the motor to recover the commanded joint position. If the resulting motion satisfies the conditions in equation (5), the MRA enters the velocity-opposing dynamic branch. Consequently, the controller does not impose a sudden disappearance of MRA torque at yielding; the actuator changes from pre-yield holding to post-yield, motion-dependent resistance.

Before toe-off, the support branch is released, the MRA is demagnetized, and the nominal calf gain is restored so that the leg can retract and reposition freely. This deployment strategy preserves the semi-active nature of the MRA: the actuator is used to resist harmful motion or share load during stance, but it is released when its resistance would interfere with swing or leg recovery.

6. Model and locomotion validation

6.1. MRA torque prediction evaluation

The trained MRA torque models are evaluated on ten unseen trajectories that are not used for model fitting. The measured joint displacement, joint velocity, and coil current are used as model inputs, and the predicted torque is compared with the measured actuator torque. Figure 3 shows representative prediction results on an unseen trajectory at the interpolation current of 1.2 A. Three representative MR-current-on sections are selected from the same motion trajectory, while the remaining parts of the trajectory are kept in the current-off state. The selected sections correspond to an early-motion section, a mid-motion section, and the full-motion section of the trajectory. This design evaluates whether each model can reproduce the MRA torque response when current is applied at different stages of the same

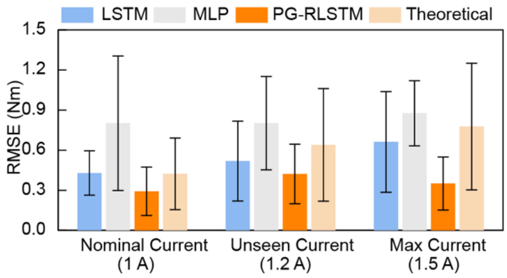

Figure 4. Torque-prediction RMSE from independent playback evaluation on ten unseen MRA trajectories for the MLP, LSTM, PG-RLSTM, and Bouc–Wen baseline models at the nominal current (1.0 A), interpolation current (1.2 A), and maximum current (1.5 A).

motion input. The recurrent models follow the measured torque trend more consistently than the feed-forward MLP, especially within the current-on sections where recent motion and current history strongly influence the MRA response.

Figure 4 reports the RMSE values over the ten unseen trajectories, with results expressed as mean $\pm$ sample standard deviation. Across all three current conditions, the PG-RLSTM achieves the lowest prediction error among the evaluated models. At the nominal current of 1.0 A, the PG-RLSTM obtains an RMSE of 0.292 ± 0.181 Nm, compared with 0.429 ± 0.166 Nm for the LSTM, 0.802 ± 0.503 Nm for the MLP, and 0.423 ± 0.268 Nm for the Bouc–Wen baseline. At the interpolation current of 1.2 A, the PG-RLSTM again gives the lowest error, 0.422 ± 0.223 Nm, while the LSTM, MLP, and Bouc–Wen baseline give 0.519 ± 0.299, 0.802 ± 0.349, and 0.640 ± 0.421 Nm, respectively. At the maximum tested current of 1.5 A, the PG-RLSTM maintains the best performance with an RMSE of 0.350 ± 0.199 Nm, compared with 0.662 ± 0.377 Nm for the LSTM, 0.876 ± 0.244 Nm for the MLP, and 0.777 ± 0.474 Nm for the Bouc–Wen baseline.

Across the three current conditions, the PG-RLSTM achieves an average RMSE of 0.355 Nm, which is lower than the LSTM (0.537 Nm), MLP (0.827 Nm), and Bouc–Wen baseline (0.613 Nm). This corresponds to average RMSE reductions of 33.9%, 57.1%, and 42.1% relative to the LSTM, MLP, and Bouc–Wen baseline, respectively. This shows that the Bouc–Wen component captures the main current-dependent hysteretic trend, while the recurrent residual branch improves prediction of transient and history-dependent torque errors. The PG-RLSTM is therefore selected as the dynamic MRA torque model for the actuator-aware locomotion simulation.

6.2. Actuator-aware locomotion simulation

The baseline Go1 policy and the MRA-SAL policy are trained and evaluated using matched observation spaces, action spaces, reward terms, terrain curricula, command sampling, policy frequency, and random-seed sets. For each policy-training condition, three independent locomotion policies are trained using different random seeds while all other training settings are held fixed. Each independently trained policy is subsequently evaluated over three rollouts. In table 2, configurations A and B share the same three conventionally trained policies, whereas configurations C and D share the same three actuator-aware policies; the paired configurations differ only in whether SAL is disabled or enabled during evaluation. For table 3, three policies are independently trained for each fixed actuator representation using the same three training seeds. For every metric, the three rollout results are first averaged within each training seed, and the final value is reported as the mean $\pm$ sample standard deviation across the three seed-level averages. Figure 5 compares the training reward curves. The MRA-SAL policy initially learns more slowly because the simulated plant includes the additional semi-active MRA branch, but it later reaches a higher reward than the baseline policy. This result indicates that the actuator-aware environment remains trainable despite the added nonlinear MRA torque contribution.

Rollout metrics are then used to evaluate locomotion stability. Across both SAL settings in table 2, the actuator-aware policy maintains a slightly higher mean CoM height than the conventional policy, indicating that actuator-aware training does not destabilize locomotion. Table 2 also isolates the contribution of actuator-aware policy training. With SAL disabled, actuator-aware training reduces the forward-velocity RMSE from 0.1062 to 0.0642 m s^{-1} and the absolute velocity bias from 0.0779 to 0.0098 m s^{-1}. With SAL enabled, the RMSE decreases from 0.1020 to 0.0708 m s^{-1}, while the absolute bias decreases from 0.0908 to 0.0150 m s^{-1}. Therefore, applying SAL to the conventionally trained

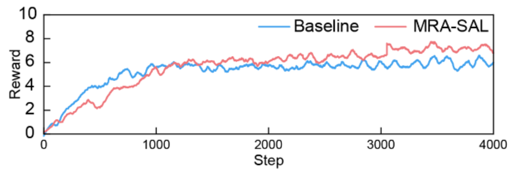

Figure 5. Training reward comparison between the baseline Go1 locomotion policy and the proposed MRA-SAL policy.

Table 2. Simulation ablation of actuator-aware policy training at $v_x^{\mathrm{cmd}} = 0.5\ \mathrm{m\,s^{-1}}$. Configuration IDs A–D are used in subsequent comparisons.

ID	Training policy	SAL	$I_{\max}$ (A)	v_x RMSE ($\mathrm{m\,s^{-1}}$)	v_x bias ($\mathrm{m\,s^{-1}}$)	Mean CoM height (m)
A	Conventional	Off	0.0	0.1062 ± 0.0077	$+0.0779 \pm 0.0122$	0.3280 ± 0.0006
B	Conventional	On	1.5	0.1020 ± 0.0149	$+0.0908 \pm 0.0136$	0.3285 ± 0.0015
C	Actuator-aware	Off	0.0	0.0642 ± 0.0057	$+0.0098 \pm 0.0110$	0.3323 ± 0.0002
D	Actuator-aware	On	1.5	0.0708 ± 0.0026	-0.0150 ± 0.0058	0.3337 ± 0.0008

Table 3. Influence of the actuator model used during policy training on robot-level locomotion performance under matched simulation conditions at $v_x^{\mathrm{cmd}} = 0.5\ \mathrm{m\,s^{-1}}$.

Training model	v_x RMSE ($\mathrm{m\,s^{-1}}$)	v_x bias ($\mathrm{m\,s^{-1}}$)	CoM-z RMSE (m)	Roll–pitch RMS (deg)	Yaw-rate RMSE ($\mathrm{rad\,s^{-1}}$)
MLP	0.0851 ± 0.0071	-0.0352 ± 0.0049	0.0143 ± 0.0004	0.7237 ± 0.3241	0.2602 ± 0.0194
Bouc–Wen	0.0713 ± 0.0075	-0.0351 ± 0.0048	0.0139 ± 0.0002	0.6706 ± 0.2919	0.2581 ± 0.0033
LSTM	0.0703 ± 0.0072	-0.0307 ± 0.0045	0.0143 ± 0.0005	0.6043 ± 0.2223	0.2524 ± 0.0052
PG-RLSTM	$\mathbf{0.0606} \pm 0.0075$	$\mathbf{-0.0265} \pm 0.0053$	$\mathbf{0.0134} \pm 0.0005$	$\mathbf{0.5177} \pm 0.3041$	$\mathbf{0.2203} \pm 0.0045$

policy at evaluation time does not reproduce the tracking improvements obtained through actuator-aware training.

The influence of the actuator model used during policy training is summarized in table 3. The MLP-based policy produces the poorest overall locomotion performance, consistent with its less accurate and more fluctuating MRA-torque predictions. These fluctuations act as artificial joint-level disturbances in the training environment and can interfere with policy learning. In contrast, the Bouc–Wen, LSTM, and PG-RLSTM representations generate smoother torque responses and produce broadly comparable locomotion behavior. This suggests that reinforcement learning can accommodate moderate differences among actuator representations, provided that their outputs remain sufficiently smooth and physically plausible. PG-RLSTM was selected because its predicted torque most closely matches the experimentally measured MRA response while maintaining smooth dynamic behavior.

6.3. Disturbance-damping evaluation in simulation and hardware

This subsection evaluates the disturbance-triggered dynamic damping branch of MRA-SAL during locomotion. The same local damping logic is used in simulation and on hardware: when a loaded calf joint moves away from its policy target, MR current is enabled and the MRA provides velocity-opposing resistance. In simulation, the dynamic damping response is evaluated by applying a 100 N vertically downward external force to the robot body during walking. For each hardware disturbance trial conducted during walking, a 10 kg mass was released from rest through a guided vertical mechanism from a nominal vertical distance of 0.30 m above the contact surface on the top of the robot body (figure 6(a)). The impact velocity was approximately 1.5 $\mathrm{m\,s^{-1}}$. After initial contact, a mechanical travel limit restricted the additional downward displacement of the mass to approximately 0.05 m, thereby limiting compression of the robot and preventing excessive loading.

Ten repeated datasets in which the impact was applied to the FR–RL diagonal stance pair were used for the reported hardware statistics. The estimated root-position and root-velocity signals available as the observations within the hardware RL control pipeline were recorded and used as proxies for vertical CoM motion. The vertical CoM displacement was calculated from the decrease in body height over the impact window, and the peak downward CoM velocity was defined as the minimum vertical velocity within the same window. The current-dependent response is evaluated at 0, 0.5, 1.0, and 1.5 A.

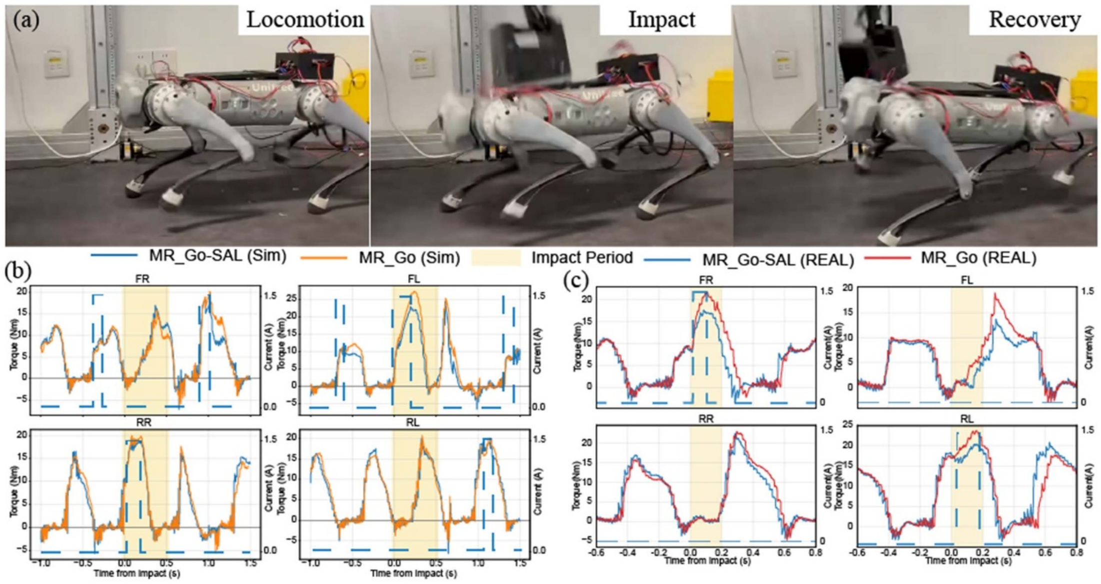

Figure 6. Representative disturbance response in simulation and hardware. The simulation result reports the PD calf-joint torque, whereas the hardware result reports the motor-side calf torque.

Figure 6(b) shows the simulated motor-side PD calf-joint torque during the 1.5 A MR-current-on disturbance interval. For the loaded FL and RR stance pair, the mean/peak torques are 12.073/23.015 Nm and 17.266/18.641 Nm, respectively. Relative to the paired no-current reference evaluated over the same time samples, the peak torque decreases by 14.5% for FL and 6.8% for RR, while the mean torque decreases by 6.7% and 2.9%, respectively. The FR and RL joints are not activated by the scheduler during this simulated impact because they are in swing, and therefore show a similar torque response between the current-on and no-current cases.

In hardware, figure 6(c) shows that the impact mainly loads the FR and RL calf joints, which are in stance during the disturbance-relevant MR-current-on intervals. For FR, MR_Go-SAL reduces the mean motor torque from 8.962 to 7.661 Nm and the peak torque from 10.775 to 8.793 Nm, corresponding to reductions of 14.5% and 18.4%, respectively. For RL, the mean motor torque decreases from 10.587 to 9.149 Nm, and the peak torque decreases from 11.840 to 10.203 Nm, corresponding to reductions of 13.6% and 13.8%. These results indicate that enabling the dynamic damping branch reduces the disturbance-related motor-side calf-torque response in hardware.

The body-level disturbance response is summarized using the minimum CoM height and the peak downward CoM velocity during impact. In simulation, increasing MR current improves both metrics: the minimum CoM height increases from 0.1851 ± 0.0011 m at 0 A to 0.2124 ± 0.0009 m at 1.5 A, while the peak downward CoM velocity becomes less negative, changing from -0.7392 ± 0.0147 to -0.5964 ± 0.0174 m s^{-1}. The hardware results show the same current-dependent attenuation trend. As shown in table 4, increasing current raises the minimum CoM height from 0.1932 ± 0.0054 m at 0 A to 0.2209 ± 0.0094 m at 1.5 A. The peak downward CoM velocity also improves from -0.6626 ± 0.1647 to -0.4279 ± 0.1008 m s^{-1}. These results indicate that higher MR current increases the available semi-active resistance and reduces impact-induced body drop.

To further evaluate motor load relief during the hardware disturbance trials, the mean calf motor torque is computed over a 0.1 s impact window, as shown in figure 7. The largest motor torque responses occur in the FR and RL calf joints. As MR current increases from 0 to 1.5 A, the FR calf motor torque decreases from 18.73 ± 1.41 to 16.03 ± 1.52 Nm, and the RL calf motor torque decreases from 16.93 ± 1.76 to 13.52 ± 1.69 Nm. Together with the improved CoM-height and downward-velocity metrics, this reduction indicates that the dynamic damping branch attenuates the disturbance while reducing the motor calf torque demand.

Beyond the current-dependent damping response, table 5 examines the additional contribution of actuator-aware training by comparing the conventional and actuator-aware policies with SAL enabled in both cases (Configurations B and D). Actuator-aware training reduces the mean absolute vertical acceleration deviation from gravity from 1.681 ± 1.668 to 0.852 ± 1.372 m s^{-2}. The mean FR and RL calf motor torques decrease from 17.51 ± 1.84 and 15.27 ± 1.23 Nm to 16.03 ± 1.52 and 13.52 ± 1.69 Nm, respectively.

Table 4. Current-dependent body-level disturbance response in simulation and hardware.

Dataset	Current (A)	Min. CoM-z (m)	Peak downward v_z (m/s^{-1})
Hardware	0.0	0.1932±0.0054	−0.6626±0.1647
Hardware	0.5	0.1970±0.0031	−0.5957±0.1673
Hardware	1.0	0.2117±0.0071	−0.4427±0.1276
Hardware	1.5	0.2209±0.0094	−0.4279±0.1008
Simulation	0.0	0.1851±0.0011	−0.7392±0.0147
Simulation	0.5	0.2030±0.0005	−0.6600±0.0046
Simulation	1.0	0.2072±0.0003	−0.6348±0.0037
Simulation	1.5	0.2124±0.0009	−0.5964±0.0174

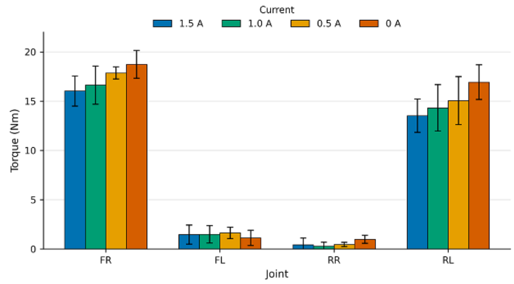

Figure 7. Hardware motor-side calf torque during the 0.1 s impact window under different MR-current settings. Bars show the mean estimated electric-motor torque for each calf joint, and error bars indicate the sample standard deviation across repeated trials.

Table 5. Configuration-level hardware disturbance comparison over the 0.1 s impact window (mean ± SD, $n = 10$). The mean absolute vertical acceleration deviation from gravity, $\overline{|a_z - g|}$, and the mean front-right (FR) and rear-left (RL) calf motor torques are reported. Configurations A–D are defined in table 2.

ID	Mean $\|a_z - g\|$ (m s^{-2})	Mean FR calf torque (Nm)	Mean RL calf torque (Nm)
A	5.216±1.474	20.84±1.29	17.98±1.05
B	1.681±1.668	17.51±1.84	15.27±1.23
C	4.893±1.219	18.73±1.41	16.93±1.76
D	0.852±1.372	16.03±1.52	13.52±1.69

Because SAL is enabled in both configurations, this comparison indicates that actuator-aware training improves the disturbance response beyond applying the scheduler to a conventionally trained policy.

6.4. Quasi-static support, calf torque relief, and COT

Figures 8(a) and (b) show the stiffness-relief mechanism during locomotion. When the front-right (FR) calf joint is in contact and its velocity is close to zero, MRA-SAL enables the quasi-static support branch and reduces the motor-side proportional gain. The MRA then supplies semi-active holding resistance, allowing the electric calf motor to maintain the stance posture with lower active torque. When the joint begins moving again, the support condition is no longer satisfied, the MRA is demagnetized and released, and the nominal motor stiffness is restored to avoid constraining leg recovery.

Figure 8(c) compares the FR hip, thigh, and calf torque distributions between the conventional Go1 and MR_Go-SAL. The hip and thigh distributions remain broadly similar, which is consistent with the hardware configuration because the MRA is installed at the calf joint rather than at the proximal joints. The clearest change appears at the calf joint: MR_Go-SAL reduces the high-torque tail and shifts the distribution toward a lower loading range. This indicates that the quasi-static support branch reduces calf motor loading without substantially redistributing the load to the hip or thigh.

Table 7 reports the mean fraction of time that the calf joints satisfy the support-activation condition. The activation time is larger at the lower walking speed and increases under the 3 kg payload, reaching 31.34% on hardware and 26.40% in simulation at 0.3 m s^{-1}. This trend is consistent with the intended

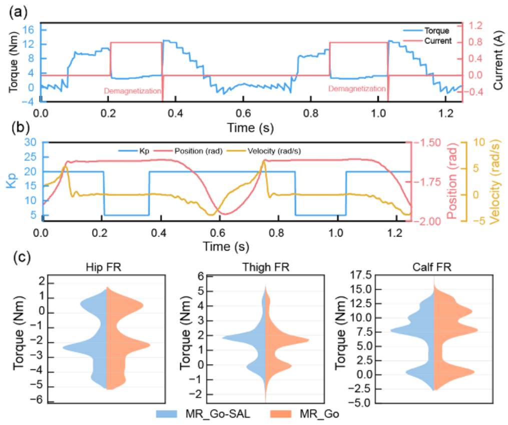

Figure 8. Quasi-static support and motor-torque relief during locomotion. (a), (b) MR_Go-SAL activates the holding branch allowing the motor-side K_p to be reduced during support. (c) Violin plots compare front-right hip, thigh, and calf torque distributions between Go1 and MR_Go-SAL; the main reduction appears at the calf joint.

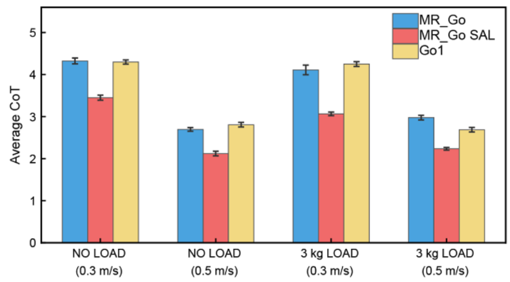

Figure 9. Average cost of transport (COT) of conventional Go1, baseline MR_Go, and MR_Go-SAL. Results are shown for no-load walking and for a 3 kg payload at commanded speeds of 0.3 and 0.5 m s^{-1}. MR_Go-SAL achieves the lowest COT in all evaluated conditions. Bars report the mean and error bars show the sample standard deviation ($n = 3$).

operating regime of the quasi-static branch: lower speed and higher load provide longer loaded stance intervals during which the MRA can share calf-joint support.

The COT is calculated from the total battery-side electrical power measured by the onboard battery management system (BMS):

$$\mathrm{COT} = \frac{\bar{P}}{mg\bar{v}} = \frac{\int_0^T P(t)\,\mathrm{d}t}{mgd}, \tag{21}$$

where $P(t) = V_{\mathrm{bat}}(t) I_{\mathrm{bat}}(t)$ is the instantaneous battery-side power recorded by the BMS, $\bar{P}$ is the mean power over the test, T is the test duration, d is the measured travel distance, and $\bar{v} = d/T$ is the measured average walking speed. The mass m denotes the total test mass, including the payload when present. The conventional Go1 has a base mass of 12.0 kg, whereas MR_Go and MR_Go-SAL have a base mass of 12.6 kg, including the MRAs and their controller. Because the measurement is taken at the battery output, it includes the complete electrical consumption of the robot and, for MR_Go and MR_Go-SAL, the MRA current controllers and excitation coils. The total mass, measured average speed, travel distance, and COT values in figure 9 are reported in table 6.

Figure 9 reports the average COT of the conventional Go1, baseline MR_Go, and MR_Go-SAL under no-load and 3 kg payload conditions at commanded walking speeds of 0.3 and 0.5 m s^{-1}. MR_Go-SAL achieves the lowest COT in all four conditions, reducing COT relative to baseline MR_Go by 20.2% at 0.3 m s^{-1} without payload, 25.4% at 0.3 m s^{-1} with the 3 kg payload, 21.3% at 0.5 m s^{-1} without payload, and 24.9% at 0.5 m s^{-1} with the 3 kg payload.

The baseline MR_Go condition provides an important control comparison because it uses the MR-integrated robot hardware without the proposed SAL scheduling strategy. Its COT is comparable to or higher than that of the conventional Go1, indicating that the efficiency improvement is not caused by the passive presence of the MR hardware alone. Instead, the reduction appears when the MRA is actively scheduled for stance support and motor-gain relief. The COT trend is also consistent with the activation

Table 6. Cost-of-transport results for the conventional Go1, baseline MR_Go, and MR_Go-SAL. All tests were conducted over a measured distance of 30 m. COT values are reported as mean $\pm$ sample standard deviation ($n = 3$).

Platform	Payload (kg)	Total mass (kg)	Mean measured speed (m/s^{-1})	COT
Go1	0	12.0	0.28	4.297 ± 0.050
Go1	0	12.0	0.49	2.807 ± 0.054
Go1	3	15.0	0.31	4.249 ± 0.058
Go1	3	15.0	0.53	2.687 ± 0.055
MR_Go	0	12.6	0.27	4.323 ± 0.070
MR_Go	0	12.6	0.48	2.695 ± 0.044
MR_Go	3	15.6	0.32	4.109 ± 0.115
MR_Go	3	15.6	0.52	2.976 ± 0.054
MR_Go-SAL	0	12.6	0.26	3.450 ± 0.060
MR_Go-SAL	0	12.6	0.48	2.121 ± 0.056
MR_Go-SAL	3	15.6	0.33	3.065 ± 0.044
MR_Go-SAL	3	15.6	0.52	2.235 ± 0.033

Table 7. Mean MRA support activation time for the four calf joints in simulation and hardware.

Hardware			Simulation		
Payload	Speed	Mean activation time	Payload	Speed	Mean activation time
0 kg	0.3 m s^{-1}	27.35%	0 kg	0.3 m s^{-1}	22.28%
0 kg	0.5 m s^{-1}	20.20%	0 kg	0.5 m s^{-1}	19.43%
3 kg	0.3 m s^{-1}	31.34%	3 kg	0.3 m s^{-1}	26.40%
3 kg	0.5 m s^{-1}	26.06%	3 kg	0.5 m s^{-1}	21.35%

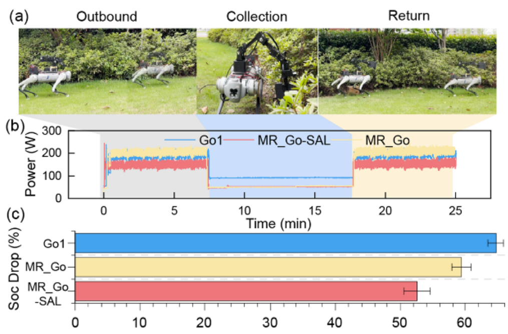

Figure 10. Field-task energy evaluation during the 300 m pickup-and-return mission. (a) Task sequence comprising a 150 m outbound traversal, item collection during a 10 min stationary period using the onboard robotic arm, and a 150 m loaded return traversal. (b) Power profiles of the conventional Go1, baseline MR_Go with SAL disabled during locomotion but conventional MR holding support retained during stationary standby, and MR_Go-SAL throughout the three task stages. (c) Battery state-of-charge (SoC) depletion over the complete mission.

times in table 7: conditions with longer support activation, especially lower-speed and payload walking, show larger COT reductions. Accordingly, the reported COT comparison accounts for both motor-system consumption and MRA excitation power.

6.5. Field task evaluation of power consumption and thermal response

To evaluate the system-level effect of MRA-SAL under a longer-duration task, MR_Go-SAL is compared with a conventional Go1 and a baseline MR_Go with the semi-active scheduler disabled in a pickup-and-return field experiment. As shown in figure 10(a), each robot first traverses a 150 m outbound segment, remains stationary at the destination for 10 min, collects the item using the onboard robotic arm, and then returns along the same 150 m track while carrying the collected item.

All three conditions use the same travel distance, payload, task sequence, and initial battery state-of-charge (SoC). The conventional Go1 uses its original actuation configuration, whereas the baseline MR_Go and MR_Go-SAL use the same MR-integrated hardware and the same battery pack. SAL is disabled during locomotion for the baseline MR_Go and enabled for MR_Go-SAL. During stationary item

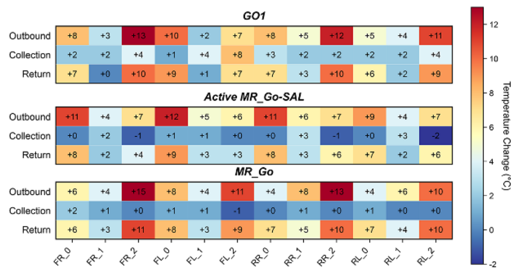

Figure 11. Temperature histories of the 12 actuators during the field experiment. The calf motors show the clearest temperature reduction because they are directly assisted by the MRA-SAL support branch.

collection, the baseline MR_Go retains the conventional predefined MR holding support used in the previous platform. This protocol allows the contribution of MRA-SAL during locomotion to be evaluated on the same physical platform. Battery SoC, electrical power, and actuator temperature are recorded throughout the trial.

Figure 10(b) shows that MR_Go-SAL consumes lower mean electrical power than both comparison conditions during the two walking stages. During outbound walking, the mean power is 143.0 W for MR_Go-SAL, compared with 198.3 W for the baseline MR_Go and 189.9 W for the conventional Go1. During item-collection standby, the baseline MR_Go retains its conventional predefined MR holding support despite SAL being disabled during locomotion. Consequently, MR_Go-SAL and the baseline MR_Go exhibit comparable mean power consumption of 51.8 and 52.2 W, respectively, both substantially lower than the 92.0 W measured for Go1. During loaded return walking, MR_Go-SAL consumes 152.4 W, compared with 208.7 W for the baseline MR_Go and 194.8 W for Go1.

As shown in figure 10(c), the SoC depletion over the complete mission is 52.5 ± 2.0 percentage points for MR_Go-SAL, 59.3±1.5 percentage points for the baseline MR_Go, and 64.7±1.2 percentage points for the conventional Go1 (mean ± SD, $n = 3$). MR_Go-SAL therefore reduces SoC depletion by 11.5% relative to the matched MR_Go baseline and by 18.9% relative to the conventional Go1. All three conditions were repeated three times under matched initial SoC, payload, speed, travel distance, and operating conditions.

Figure 11 compares the actuator-temperature changes of the conventional Go1, the baseline MR_Go with SAL disabled, and MR_Go-SAL. The clearest difference occurs at the calf motors, which are directly assisted by the MRAs. Across the outbound and return walking stages, the mean temperature rise of the four calf motors is 19.75 °C for Go1, 22.25 °C for the baseline MR_Go, and 11.50 °C for MR_Go-SAL. MR_Go-SAL therefore reduces the mean calf-motor temperature rise by 48.3% relative to the matched MR_Go baseline and by 41.8% relative to Go1. During item-collection standby, the mean calf-temperature changes are +4.50 °C, −0.25 °C, and −1.00 °C for Go1, MR_Go, and MR_Go-SAL, respectively. Across all 12 motors over the complete field task, the corresponding mean temperature rises are 16.25 °C, 15.17 °C, and 13.00 °C. These results show that the principal thermal benefit of MRA-SAL is concentrated at the directly assisted calf motors, while the proximal-joint temperature changes exhibit a less direct trend.

7. Discussion

The results indicate that the benefit of MRA-SAL comes from combining physics-guided MRA modeling with mode-dependent semi-active scheduling. At the actuator level, the Bouc–Wen-informed PG-RLSTM improves torque prediction by combining a current-dependent hysteresis baseline with a recurrent residual term. This is important for MR actuation because the torque response is not determined by current alone, but also by recent motion history, passive drag, and trajectory-dependent nonlinearities.

Embedding the learned MRA model into the Isaac Gym low-level torque loop also reduces the mismatch between the simulated and physical semi-active system. Instead of treating the MRA as an unmodeled deployment-time correction, the actuator-aware simulation exposes the locomotion policy to the additional MR resistance during rollout. This is particularly relevant because the policy is updated at 50 Hz while the low-level torque loop runs at 200 Hz; unmodeled damping at the calf joint could otherwise alter stance timing, joint excursion, and corrective actions on hardware.

The distinction between dynamic damping and quasi-static holding is central to using the same MRA during locomotion. The dynamic branch is useful when a loaded calf joint moves away from its target, where velocity-opposing MR resistance can attenuate disturbance-induced motion. The quasi-static branch is useful when the calf joint is loaded and nearly stationary, where pre-yield MR resistance can share the stance load with the electric motor. This mode separation explains why MRA-SAL can reduce calf motor loading and improve disturbance response without continuously resisting swing, retraction, or leg recovery.

Accordingly, the two branches are evaluated using operating-regime-specific metrics: CoM and motor-torque responses characterize dynamic damping, whereas motor-load and energy-consumption results characterize quasi-static holding. In the holding regime, the energy-saving benefit arises from MRA-supported gain reduction rather than gain reduction alone: retaining the nominal K_p leaves the motor responsible for nearly the same stance effort, whereas reducing K_p without MRA support risks increased tracking error and degraded posture regulation. The complete scheduler selects the appropriate branch according to the gait and disturbance state. Improved dynamic torque prediction mainly supports actuator-aware policy training, whereas the energy-saving and disturbance-attenuation benefits arise from the scheduled operation of the physical MRA.

The selected low-gain setting in the quasi-static branch should be interpreted as a practical stability compromise. In principle, the largest motor-torque relief would be obtained by setting the effective calf proportional gain to zero and allowing the MRA to provide all stance holding resistance. In hardware tests, however, fully removing calf stiffness made the contact response sensitive to modeling error, foot-force variation, and timing mismatch between MR current buildup and leg loading. Therefore, a nonzero low-gain setting, $K_{p,i}^{\mathrm{eff}} = 5$, was used instead of $K_{p,i}^{\mathrm{eff}} = 0$. This retains a small amount of active motor stiffness for posture regulation while still allowing the MRA to offload part of the stance support demand.

The 1.5 Hz gait frequency used for policy training follows the same practical reasoning. The target operating regime in this study is low-speed locomotion up to approximately 0.5 $\mathrm{m\,s^{-1}}$. At this speed range, the stance intervals are long enough for the quasi-static MRA branch to contribute holding resistance when the calf joint velocity is near zero. With the fixed 1.5 Hz gait, the maximum reliably achievable speed was approximately 0.6 $\mathrm{m\,s^{-1}}$. Commands of 0.8–1.0 $\mathrm{m\,s^{-1}}$ would require stride lengths beyond the practical leg workspace and would primarily produce trajectory-tracking errors. Evaluation at these speeds therefore requires policy retraining at a higher gait frequency and corresponding retuning of the MRA scheduler. A higher gait frequency would shorten both the stance-support interval and the disturbance-damping activation window.

Beyond gait speed, the tested payload was limited to 3 kg to remain within the reliable load-carrying envelope of the small quadruped and to obtain repeatable comparative data. Exceeding the load-carrying capability of the platform can prevent the joints from lifting the body or accurately tracking the commanded leg trajectories, resulting in large joint errors, actuator saturation, and possible gait failure. Such errors may also cause frequent or persistent activation of the dynamic damping gate, thereby confounding the evaluation of the scheduler. The resulting performance would therefore be dominated by the motor and kinematic limits of the platform rather than by the effectiveness of the MRA-SAL framework.

The disturbance tests focused on vertical loading, which directly excites the MRA-equipped calf joints. Forward–backward disturbances may also excite these sagittal-plane joints, whereas lateral disturbances are more likely to load the hip abduction/adduction joints, where MRAs are not installed. The present results should therefore not be interpreted as direction-independent disturbance rejection; stronger and multidirectional disturbances will be evaluated in future work.

The simulation and hardware disturbance results support the same qualitative mechanism: increasing current increases the available semi-active resistance and reduces body drop during impact. However, these results should not be interpreted as proof of globally optimal current scheduling. The present high-current event-triggered strategy is a practical choice for short and unpredictable disturbances. Future work should investigate adaptive current scheduling that balances damping benefit, coil energy consumption, actuator temperature, and long-term material constraints.

Several limitations remain. First, the learned torque model is trained within a bounded operating range of current, displacement, velocity, and gait condition, and additional data may be required for aggressive locomotion or higher-speed impacts. Second, long-term effects such as MR fluid temperature variation, hysteresis drift, particle sedimentation, and actuator aging are not explicitly modeled. Third, the present validation focuses on one MRA-integrated quadruped platform, so the transferability of the modeling and scheduling framework to other robot morphologies and MR actuator designs should be investigated in future work.

8. Conclusion

This paper presented MRA-SAL, a physics-guided modeling and semi-active control framework for quadruped locomotion with calf-joint MRAs. To represent the nonlinear and history-dependent MRA response, a current-aware Bouc–Wen model was combined with a PG-RLSTM residual model. The resulting torque predictor achieved the lowest RMSE among the evaluated models, with an average RMSE of 0.355 Nm across the tested current conditions, demonstrating that the Bouc–Wen baseline and recurrent residual term can jointly capture the main current-dependent and transient MRA torque behavior.

The learned MRA model was embedded into the Isaac Gym low-level torque loop to form an actuator-aware locomotion simulation. Based on this model, the MRA-SAL scheduler separated two physically different semi-active functions: quasi-static stance support during loaded low-velocity contact and velocity-opposing dynamic damping during disturbance-induced target-escaping motion. Simulation results showed current-dependent disturbance attenuation, while hardware disturbance experiments confirmed the same qualitative trend through increased minimum CoM height and reduced peak downward velocity under higher MR current.

Real-robot locomotion experiments further showed that MR_Go-SAL reduces calf motor torque loading, lowers COT relative to the baseline MR_Go condition, and decreases electrical power consumption and calf-motor temperature rise during a field pickup-and-return task. These results indicate that magnetorheological actuation can provide practical locomotion benefits when its current-dependent material behavior is explicitly modeled and scheduled according to the locomotion state. Future work will extend the model to include thermal variation, hysteresis drift, and long-term actuator aging, and will develop adaptive current scheduling that balances damping performance, coil energy consumption, and smart-material durability.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No. U21A20119, 52105081 and 62303002), State Key Laboratory of Micro-Spacecraft Rapid Design and Intelligent Cluster (Grant No. MS01240110), Major Project of Anhui Province's Science and Technology Innovation Breakthrough Plan (202423h08050003), Anhui's Key R&D Program of China (Grant No. 202104a05020009), USTC start-up funding (Grant No. KY2090000067), Key Program of Natural Science Foundation of Anhui Higher Education Institutions of China (Grant No. 2022AH050068), and the Fundamental Research Funds for the Central Universities (SC5290005162).

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Supplementary Video 1 available at: https://doi.org/10.1088/1361-665X/ae9e72/data1.

Conflict of interest

None of the authors has a conflict of interest to disclose.

ORCID iDs

Wenbo Li 0009-0001-7945-4958
Shuaishuai Sun 0000-0002-8695-9217

References

[1] Plante J S, St-Jean A and Lucking Bigué J P 2025 *Sci. Rep.* **15** 14363
[2] Véronneau C, Denis J, Lhommeau P, St-Jean A, Girard A, Plante J S and Bigué J P L 2022 *IEEE Robot. Autom. Lett.* **8** 896–903
[3] Pisetskiy S and Kermani M 2021 *J. Intell. Mater. Syst. Struct.* **32** 2582–600
[4] St-Jean A, Dorval F, Plante J S and Lussier-Desbiens A 2023 *IEEE Trans. Robot.* **40** 1261–72

[5] Lakatos D, Ploeger K, Loeffl F, Seidel D, Schmidt F, Gumpert T, John F, Bertram T and Albu-Schäffer A 2018 *IEEE Robot. Autom. Lett.* **3** 3908–15
[6] Ding J, Atanassov V, Panichi E, Kober J and Della Santina C 2024 *IEEE Trans. Robot.* **40** 3212–31
[7] Ding J, van Löben Sels M A, Angelini F, Kober J and Della Santina C 2023 *IEEE Robot. Autom. Lett.* **9** 255–62
[8] Hwangbo J, Lee J, Dosovitskiy A, Bellicoso D, Tsounis V, Koltun V and Hutter M 2019 *Sci. Robot.* **4** eaau5872
[9] Raffin A, Seidel D, Kober J, Albu-Schäffer A, Silvério J and Stulp F 2022 arXiv:2209.07171
[10] Bjelonic F, Lee J, Arm P, Sako D, Tateo D, Peters J and Hutter M 2023 *IEEE Robot. Autom. Lett.* **8** 1611–8
[11] Lee M, Kim H, Kim J T, Park S, Lee J H, Cho J and Hwangbo J 2025 *IEEE Robot. Autom. Lett.* **10** 12421–28
[12] Yuryev V and Hughes J 2026 arXiv:2603.04351
[13] Spencer J B Dyke S, Sain M and Carlson J 1997 *J. Eng. Mech.* **123** 230–8
[14] Abdelhamed M, Ata W, Salem A and Khafagy A 2020 Effect of modified Bouc-Wen model parameters on the dynamic hysteresis of magnetorheological dampers *IOP Conf. Ser.: Mater. Sci. Eng.* **973** 012020
[15] Liu Z, Wang S, Cao Y, Cheng H and Xie L 2022 *IEEE Robot. Autom. Lett.* **7** 11609–16
[16] Wu Y, Sicard B, Kosierb P and Gadsden S A 2025 magnetorheological dampers via physics-informed neural networks *Proc. IEMTRONICS 2024: Int. IoT, Electronics and Mechatronics Conf.* vol 1 (Springer) p 165
[17] Wu Y, Sicard B, Kosierb P, Appuhamy R, Leroux A M and Gadsden S A 2024 *Physics-Informed Neural Networks for Modeling Hysteretic Behavior in Magnetorheological Dampers* ed P J Schwartz, M E Hohil and B Jensen (SPIE) p 54
[18] Li Z, Peng L, Hua D, Shen Y, Liu X, Mu H, Wu J and Zhang T 2026 *Smart Mater. Struct.* **35** 065048
[19] Chun1a S, Javidan M M, Xiang3b Y and Kim J 2026 *Struct. Eng. Mech.* **98** 577–97
[20] Ferhath A A and Kasi K 2025 *Iran. J. Sci. Technol. Trans. Mech. Eng.* **49** 2387–405
[21] Lu H, Peng L, Hua D, Liu X and Yao R 2024 *Smart Mater. Struct.* **33** 085015
[22] Jiang M, Rui X, Yang F, Wang P, Li H and Wang W 2026 *J. Vib. Control* **32** 1838–51
[23] Wei S, Wang J and Ou J 2021 *Mech. Syst. Signal Process.* **149** 107316
[24] Yew E A T, Tao Z, Cai H, Kong Z, Zhou B, Zhang L and Sun S 2026 *IEEE Robot. Autom. Lett.* **11** 5670–77
[25] Ashtiani M, Hashemabadi S and Ghaffari A 2015 *J. Magn. Magn. Mater.* **374** 716–30
[26] De Vicente J, Klingenberg D J and Hidalgo-Alvarez R 2011 *Soft Matter* **7** 3701–10
[27] Yew E A T, Li W, Mai X, Chen Z, Zhang L and Sun S 2026 *Smart Mater. Struct.* **35** 085003
[28] Ma H, Chen B, Qin L and Liao W H 2017 *Smart Mater. Struct.* **26** 035013