BRNO UNIVERSITY OF TECHNOLOGY FACULTY OF INFORMATION TECHNOLOGY

Excel@FIT 2026
http://excel.fit.vutbr.cz

AI-Based Control of an e-Bike Rear MR Damper

Bc. Pavel Kejík*

Abstract
Rear suspension on electric mountain bikes must balance comfort over terrain inputs with support under rider-induced loads such as pedalling and pumping. This work addresses the problem using a semi-active magnetorheological damper. Classical laws such as Skyhook fix this trade-off at design time, which limits adaptability. This work presents a compact GRU controller (7 inputs $\rightarrow$ 24 hidden $\rightarrow$ 1 output) trained on a closed-loop simulation of a magnetorheological damper coupled to a three-mass suspension model. Training uses Evolution Strategies on mixed road, obstacle, and rider-event scenarios, with a user preference slider controlling the comfort–support trade-off. On a mixed simulated validation set, the controller improves frequency-weighted RMS acceleration by about 8 % over Skyhook in comfort mode, with gains of about 11 % on forest trail and 10–12 % on rough trail surfaces. On-bike logs confirm lower weighted RMS and reduced peak stroke on rider-induced events, and on forest trail show about 8.4 % lower weighted RMS and about 33 % lower average current in comfort mode. The controller runs at 1 kHz on an ESP32-S3, with GRU inference taking about 114 µs per step.

*xkejik06@stud.fit.vutbr.cz, *Faculty of Information Technology, Brno University of Technology*

1. Introduction

During a single ride, the rear suspension must handle seated climbing, sprinting, discrete impacts, and rough descents. Semi-active MR dampers address this by modulating damper force through coil current on a millisecond time scale [1].

The Skyhook law [2] maps sprung-mass velocity to damping force. It is simple and robust, but it does not distinguish terrain-driven motion from rider-driven motion. The comfort–support trade-off is therefore fixed by tuning and cannot be adjusted by the rider without recalibration.

Prior work on semi-active suspension control includes lookup-based Skyhook variants, clipped hybrid laws, and MPC. In this application, MPC can be computationally demanding at 1 kHz on a battery-powered MCU. Reinforcement-learning-based controllers have also been reported, but direct deployment on low-cost embedded hardware remains difficult.

This work uses a compact recurrent controller (GRU with 24 hidden units) trained on a closed-loop simulation. A gradient-free Evolution Strategies (ES) optimiser is used, and the controller takes seven engineered features together with a user preference slider $s \in [0, 1]$. The output is the damper current command. Fig. 1 shows the hardware setup, Fig. 2 the online control path, and Fig. 3 the offline training pipeline. The system is built around a custom MR damper developed at the Faculty of Mechanical Engineering, Brno University of Technology, with a response time of approximately 2 ms.

The key contribution is a compact 24-unit GRU controller that outperforms a well-tuned Skyhook baseline in simulation, together with a training framework based on a continuous comfort–support slider and a real-time 1 kHz ESP32-S3 implementation validated on bike logs.

2. Online Control Pipeline

The online path (Fig. 2) runs as a fixed 1 kHz loop. Two IMUs provide vertical accelerations on the sprung mass (frame) and the unsprung mass (swingarm+wheel); a gyroscope provides pitch rate; and a magnetic stroke sensor measures relative rear-suspension displacement. Signal conditioning combines Mahony attitude estimation [3], band-pass filtering, and leaky integrators with a zero-velocity update (ZUPT) gate to estimate body velocity.

The GRU takes seven inputs – two band-passed accelerations, relative stroke velocity, estimated body velocity, previous current command, stroke, and slider – and outputs the next current setpoint I_{ref}. The damper driver tracks I_{ref} within the coil's 0–1.7 A range.

Measured inference time for the GRU on the ESP32-S3 is about 114 µs. The inference kernel is custom-optimised for the ESP32-S3 using SIMD instructions and currently runs in float32 precision. Quantisation was not required, since the float32 implementation already meets the timing budget and larger GRU variants did not improve the measured metrics. The slider is received over BLE via a mobile app and enters the policy directly as an input.

3. Training in Simulation

Training is performed off-line on a custom scenario library (Fig. 3). The plant is a three-mass vertical model covering the unsprung mass (swingarm with wheel), the sprung mass (frame), and the rider as a lumped mass above the frame. The MR damper is modelled by a measured $F(I, v)$ map and an approximately 2 ms current response (shown in Fig. 3).

Training uses stochastic ISO 8608 roads [4] of classes A–D, discrete obstacles (bumps, steps, drops), and textured sequences such as gravel or cobblestone. All surfaces are passed through a rolling geometric tire envelope of radius $R = 0.33$ m, which removes sub-wheel features the tire cannot physically excite. Rider-induced loads are modelled separately as forces applied to the sprung mass: a cyclic pedalling load with drifting cadence and amplitude, and smooth push pulses representing preloads, weight transfers, and bunny-hops. Scenario seeds, speed, event timing, and selected plant parameters are randomized in both training and validation.

Optimisation uses an Evolution Strategies (ES) algorithm with antithetic sampling [5]. Each candidate is evaluated on a mixed batch covering the slider values used by the rider ($s \in \{0,\ 0.5,\ 1\}$). The objective is a weighted combination of ISO 2631 weighted RMS acceleration and VDV [6], tire-load stability (DLC, contact loss), anti-bob, rider sink, stroke usage, current usage and current-switching smoothness ΔI. All metrics are normalised against a Skyhook baseline evaluated on the same scenario set.

The slider enters the system in two places. During training, it changes the objective weights towards comfort ($s = 0$) or support ($s = 1$). During inference, it is used as the 7th GRU input. This allows the same controller to cover different comfort–support settings without switching between separately tuned control laws.

4. Results

Simulation — overall. Fig. 4 reports percent improvement over Skyhook on a held-out simulated validation set at the comfort and support slider ends. At $s = 0$, weighted RMS acceleration improves by 8.1 % and VDV by 9.9 %. Average coil current is reduced by about 20 %, while tire-load stability remains essentially unchanged.

Simulation — terrain and slider. The 8 % figure is a mixed-scenario mean. On rough trail surfaces the RMS reduction grows to 10–12 %. In simulation, the forest trail theme improves by about 10.6 % at $s = 0$. Fig. 6 shows the slider sweep over $\{0, 0.25, 0.5, 0.75, 1\}$; comfort metrics decrease towards support, while tire-load stability and anti-bob improve. Intermediate slider values produce intermediate behaviour.

On-bike measurements. Fig. 5 shows a measured push event. On repeated push sequences, weighted RMS drops by about 10 % and peak stroke from 17.9 mm to 13.8 mm at $s = 0$. On single pushes, weighted RMS drops by about 18 % and peak stroke by about 10 % at $s = 0$. On forest trail, weighted RMS drops by about 8.4 % at $s = 0$, while average current is about 33 % lower.

5. Conclusions

A 24-unit GRU controller trained by Evolution Strategies on a closed-loop simulation outperforms a well-tuned Skyhook baseline within the compute budget of a low-cost MCU. The gains are small on smooth terrain, about 11 % on forest trail and 10–12 % on rough trails in simulation, and clearly visible in on-bike logs, where weighted RMS and peak stroke are reduced on rider-induced events. The same controller works across the full slider range without retraining.

Future work includes a move from the current vertical three-mass model to a half-bike model with coupled front and rear suspensions, an additional IMU on the front wheel to provide a short preview of the upcoming terrain, and a controlled user study on the bike.

Acknowledgements

This work was developed in collaboration with **GRIXI s.r.o.**, who provided the bike platform with the MR damper. The MR damper was developed at the Faculty of Mechanical Engineering, Brno University of Technology, and the control unit hardware was designed by **Ing. Václav Šimek**.

References

[1] B. F. Spencer, S. J. Dyke, M. K. Sain, and J. D. Carlson. Phenomenological model for magnetorheo-

logical dampers. *Journal of Engineering Mechanics*, 123(3):230–238, 1997. https://doi.org/10.1061/(ASCE)0733-9399(1997)123:3(230).

[2] D. Karnopp, M. J. Crosby, and R. A. Harwood. Vibration control using semi-active force generators. *Journal of Engineering for Industry*, 96(2):619–626, 1974. https://doi.org/10.1115/1.3438373.

[3] Robert Mahony, Tarek Hamel, and Jean-Michel Pflimlin. Nonlinear complementary filters on the special orthogonal group. *IEEE Transactions on Automatic Control*, 53(5):1203–1218, 2008. https://doi.org/10.1109/TAC.2008.923738.

[4] International Organization for Standardization. Mechanical vibration – road surface profiles – reporting of measured data. ISO 8608:2016, 2016.

[5] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. arXiv preprint arXiv:1703.03864, 2017. https://arxiv.org/abs/1703.03864.

[6] International Organization for Standardization. Mechanical vibration and shock – evaluation of human exposure to whole-body vibration – part 1: General requirements. ISO 2631-1:1997, 1997.