




Universal differential equations for modelling degradation of suspension dampers

Lorenz Ott^{a, b, *} , Silke Redecker^{c, b}, Torben Gräber^b, Michael Unterreiner^b, Johannes Edelmann^a, Manfred Plöchl^a

^a Institute of Mechanics and Mechatronics, TU Wien, Vienna, Austria

^b CARIAD SE Pre-Development, Stuttgart/Mönsheim, Germany

^c KIT Karlsruhe, Karlsruhe, Germany

ARTICLE INFO

Keywords:

Universal differential equations
Scientific machine learning
Suspension dampers
Physics-informed machine learning
Degradation

ABSTRACT

Suspension dampers degrade over time. They operate under a wide range of load conditions and environmental influences, causing changes in damping characteristics that may affect vehicle safety and ride comfort. Accurate simulation models that capture these changing characteristics are essential for degradation effect analysis, yet corresponding research remains limited in the literature.

This article proposes a simulation model for degraded suspension dampers based on the Universal Differential Equations framework. The modelling approach starts with an equivalent mechanical model, covering known physical effects, which is then enhanced through the integration of Neural Networks into its system dynamics. Physical consistency constraints are enforced by auxiliary losses throughout training.

Test bench measurements of functional and degraded dampers reveal that oil loss introduces strongly non-linear, transient and asymmetric changes to damper dynamics. These effects become increasingly pronounced at higher excitation frequencies and smaller stroke amplitudes. The newly developed Neural Equivalent Mechanical Model was trained and validated using the test bench data and was shown to effectively capture the dynamics induced by degradation. Further interpretation of the learned neural functions reveals insights into how oil loss influences the dynamics captured by the developed damper model.

1. Introduction

The fundamental objectives of a vehicle suspension are to isolate the vehicle body from road-induced vibrations and to improve vehicle handling through the use of spring and damper elements [1]. These objectives are achieved by tuning the suspension design parameters [2]. However, suspension components degrade over time since they operate under certain stress or load in the real environment, often involving randomness. Reports show that this ultimately results in every hundredth vehicle failing a regular check after seven years, due to faults related to suspension springs and dampers [3]. As a result of degradation, component characteristics change and affect vehicle safety and ride comfort [4–6].

To analyse the effects of degradation on components, test bench measurements of degraded components can be performed. However, in the context of effect analysis on damper and vehicle levels, the well-known adage “a picture is worth a thousand words” might be rephrased as

“a (simulation) model is worth a thousand measurements”. In addition to enabling effect analysis, damper models, if accurate, can provide a detailed understanding of changing characteristics and thus allow, for instance, conclusions on vehicle design. Consequently, there is a demand for suspension damper models that can accurately capture damping dynamics across a broad range of operating conditions and damper conditions influenced by degradation.

This work focuses on finding a model for degraded dampers, whose structure and optimised parameters are, to some degree, interpretable, which is crucial for improving knowledge based on the parametrised model.

To analyse the effect of degradation on damper characteristics, various real-world measurements of a hydraulic vehicle suspension damper with and without degradation were performed on a test bench. To replicate damper degradations for analysis and model parametrisation, the present work focuses on oil loss as the primary degradation effect, using controlled oil reduction. The choice of oil loss as a degradation

* Corresponding author at: Institute of Mechanics and Mechatronics, TU Wien, Vienna, Austria.

Email address: lorenz.ott@cariad.technology (L. Ott).

effect is argued based on existing studies: Hu et al. [7] investigated the underlying degradation mechanisms of hydraulic dampers and found that cyclic piston motion generates a high-pressure, high-temperature environment inside the damper cylinder, leading to seal cracking, resulting in oil leakage, and air ingress. Moreover, Schramm et al. [8] showed experimentally that oil loss has a more dominant effect on the change of damper force characteristics than internal pressure loss alone. Also in other studies, oil loss has been widely used to replicate damper degradation experimentally across automotive and structural applications [4,7–9].

The conducted measurements show that degradation affects the damper dynamics such that the relationship between damper states and the measured force becomes more nonlinear and transient. A change in damper dynamics was confirmed by [4,8], which presents static degradation models that do not include transient effects. Moreover, the model of [8] focuses on a rule-based logic to distinguish between various damper states rather than a model applicable to broader operating regions. In addition to vehicle suspension dampers, the influence of oil leakage on the performance of structural dampers was, e.g., analysed in [7,10] for railway vehicles.

For non-degraded dampers, the challenge of mathematically representing nonlinear, time-varying behaviour has led to substantial academic research, mainly divided into two categories: Hydraulic Models (HMs) based on first-principles from physics and Equivalent Mechanical Models (EMMs).

HMs describe damper dynamics by formulating mathematical equations for fluid-flow, fluid-and-gas compression, and the damper's geometric specifications [11]. The literature covers a broad range of HMs with varying degrees of model fidelity [12–14]. HMs are particularly useful because they reveal how changes in physical properties, such as geometric details, affect damper behaviour, enabling detailed parameter studies [13]. However, these models, while insightful, have their limitations. The accuracy of any HM relies on precise parameterisation, and obtaining accurate values typically requires a series of intricate, time-consuming characterisation experiments. Moreover, parameter dependency on operating conditions further complicates the parameterisation process. Consequently, extending an HM to cover a wide range of degradation conditions becomes impractical, since it would demand extensive experiments that are difficult, costly, and time-consuming to perform.

An EMM, on the other hand, is a model that is computationally efficient and based on relatively few independent variables to obtain a good fit to measured data. In the simplest case, the EMM consists of a damping coefficient that is multiplied by the damper velocity to approximate the damping-force. However, a linear damping coefficient is insufficient for capturing the nonlinear characteristics of dampers. This has led to the adoption of a map-based approach, where operating conditions (such as damper velocity and stroke amplitude) are varied in test bench measurements and the damping-force is measured at quasi-static velocity points [15]. Based on these measurements, a damping map is derived. Thus, the damping map exhibits locally Linear Time-Invariant (LTI) behaviour in each operating region but is globally nonlinear. Measurement results show that the assumption of local LTI behaviour becomes invalid in operating regions characterised by high frequencies, where the damper exhibits significant hysteresis [12].

To account for hysteresis, Unger et al. [1] proposed a mechanical equivalent model that reduces HM submodels to mechanical elements and integrates them into an overall model. This approach maintains low complexity relative to HMs while preserving interpretability, since spring and damper coefficients retain physical meaning. The dependence of model parameters on damper conditions is captured via a map-based approach using discrete data points fitted to test-bench experiments. However, such models are constrained by their limited structural complexity (kept low to reduce the amount of parameterisation measurements) and by the way of integrating measurements in the parameterisation process, generally referred to as “model fitting”, which

is inherently limited by local identification. These limitations can restrict the estimation accuracy of EMMs across a wide range of operating and system conditions.

Alternatively, standard data-driven models such as Neural Networks (NNs) can be trained over a set of training instances, which can then be used to produce estimates of the resulting damper dynamics [16]. While these models are able to model complex behaviour and thus potentially the effects of degradation, limited conclusions can be drawn from the trained model, as the parameters are not interpretable. Moreover, these models typically require large amounts of training data and may fail to generalise to unseen operational conditions.

The above discussion highlights the potential benefits of combining physics-based and data-driven models for simulation purposes. Hybrid simulation models should possess traits (A1–A3) that enable them to make accurate estimations of damper forces across a wide range of operating and system conditions:

- (A1) Combine known dynamics from mechanistic models with data-driven methods, preserving the identified structure of the mechanistic submodels and their interrelations. Leverage the flexibility of data-driven approaches to discover additional nonlinearities influencing the system dynamics.
- (A2) Retain model interpretability so that conclusions can be drawn from the parameterised models, and consistency with physics-based knowledge can be assessed and enforced. Moreover, the interpretability can be used to analyse and avoid stability issues.
- (A3) Utilise big datasets in a computationally efficient manner to achieve accurate models across a wide range of damper states and degradation conditions.

One promising approach to fulfil these traits is the methodology of Universal Differential Equations (UDE) [17], from scientific machine learning. “Scientific machine learning” seeks to combine machine learning methods with traditional engineering-related methods in order to utilise domain knowledge and known physical information in the learning process. The UDE approach, specifically, embeds neural networks directly into the differential equations of a mechanistic model, so that the physics-based structure is preserved and the networks learn the residual dynamics not captured by the mechanistic model.

Published methods related to advancing data-driven models with physical knowledge, such as Physics-informed Neural Networks (PINNs) [18], aim to achieve consistency with physics through a physics-constrained loss function. Li et al. [19] demonstrate accuracy gains for viscoelastic dampers by embedding constraints derived from a molecular chain model into the training loss. Li et al. [20] confirm the same benefit for viscoelastic materials under thermal-oxidative ageing, where physical constraints maintain high accuracy even with moderate training data and across varying degradation levels. In the automotive domain, a related application is the Damper-B-PINN framework [21], which uses suspension damper characteristics as physical guidance for vehicle state estimation. However, in all these approaches, the neural network remains uninterpretable.

Unlike PINNs and related concepts, the UDE framework [17] embeds neural networks directly into the differential equations of a mechanistic model. This enables a fully differentiable model, allowing gradients to be backpropagated from high-level states to any part of the computation graph. This approach has been successfully implemented in vehicle dynamics modelling [22] and hysteresis modelling [23]. However, standard UDE formulations typically learn additive residuals to the state dynamics, which limits the physical interpretability of the learned components. The proposed method builds upon the concept of embedding neural networks in the system dynamics of a mechanistic model. A complementary strategy to UDEs and similar methods for advancing physics-based models with data-driven modelling capacity is the use of residual-models [10,24–26], which learn deterministic residuals to

the outputs of physics-based models without utilising any physical knowledge during the learning process itself.

In this work, the UDE framework is applied to create a novel suspension damper model called Neural-EMM (Neural Equivalent Mechanical Model).

This work's main contributions are:

- Collection of test bench measurements with functional and degraded suspension dampers at multiple discrete degradation levels, providing a systematic dataset for model development and validation.
- Design of a task-specific Neural-EMM based on the UDE framework, embedding neural networks as multiplicative scalars on physically meaningful EMM parameters to improve modelling capacity while preserving physical interpretability.
- Design of physics-based auxiliary losses that enforce physical consistency constraints on the learned neural functions throughout training.
- Interpretation of the learned model and abstraction of physically meaningful insights into the effect of oil-loss degradation on damper characteristics.

The remainder of this paper is organised as follows. Section 2 describes the UDE framework, including its formulation, components, and scope. Section 3 presents the experimental design, data description and analysis. Section 4 discusses the detailed design of the EMM, and in Section 5, implementation details of the suggested UDE model are provided. Section 6 presents the performance evaluation results, and a short summary with conclusions is provided in Section 7.

2. Universal differential equations

2.1. General UDE form

Physical systems, such as suspension dampers [1], can be modelled as dynamical systems comprising state and output equations:

$$\frac{dx}{dt} = f(x, u, t; \theta_{ph}), \quad (1)$$

$$y(t) = g(x, u, t; \theta_{ph}), \quad (2)$$

with $x \in \mathbb{R}^n$ representing the system state vector, $u \in \mathbb{R}^k$ denoting external inputs to the system, $y \in \mathbb{R}^m$ indicating the primary variables of interest, and $\theta_{ph} \in \mathbb{R}^p$ containing the model parameters. The functions f and g are derived from physics-based knowledge, also known as “mechanistic models” [27].

Such damper models have few, physically meaningful parameters, such as stiffness and damping parameters, which makes them interpretable. The parameters can directly be measured or estimated from limited dynamometer data [1]. For models that generalise well across varying damper operating conditions, such as varying velocities, the functions f and g in the mechanistic model must accurately incorporate the physical mechanisms relevant to those operating conditions.

However, the accuracy of mechanistic models is limited by the modelled mechanisms. Identifying all relevant mechanisms poses a significant challenge. For example, in a mechanistic representation of damper elasticity, it is assumed that fluid compression is negligible [11]. However, as dampers degrade and lose oil, the remaining fluid can foam, causing a significant reduction in stiffness. In such cases, the original model performs poorly at higher degradation levels.

The main idea of a UDE approach is to embed NNs as universal approximators (UAs) inside the differential equations of the mechanistic model, preserving the base structure while allowing additional flexibility. Hereby, model accuracy can be improved by enhancing the functions f and g of the mechanistic equations.

This enhancement can be achieved in different ways: one design choice is to train NNs that introduce relevant nonlinearities into the system dynamics by estimating model parameters as functions of the

state [27]. Another approach is to introduce NNs that estimate a state residual, adjusting for deterministic discrepancies not accounted for in the original model formulation [28]. By using comprehensive data, spanning the desired operating conditions, to train the NNs, the UDE model approach can account for the relevant dynamics and nonlinearities not represented in the baseline mechanistic model.

Formally, the UDE model is described as:

$$\frac{dx}{dt} = f(x, u, \text{NN}(x, u; \theta_{NN}); \theta_{ph}), \quad (3)$$

$$y(t) = g(x, u, \text{NN}(x, u; \theta_{NN}); \theta_{ph}), \quad (4)$$

where typically feedforward NNs are used in the UDE formulation.

Two key advantages of this approach are: first, as the UDE structure already captures the mechanistic dynamics, the NN can be much simpler (fewer parameters) and therefore typically requires less data than traditional machine learning (ML) methods. Second, integrating NNs enhances the accuracy of the mechanistic model while, if chosen appropriately, preserving the overall interpretability.

2.2. UDEs in damper modelling

In Fig. 1, the learning setup of a UDE model is illustrated for damper modelling. The top-left panel shows the measured behaviour; in this case, the damper force F_{meas} . These measurement data are the basis for developing a physics-based model that approximates the observed behaviour with F_{ph} and corresponding model equations f and g , shown in the top-middle panel. With the measurement data and the physics-based model, shortcomings of the physical model can be identified. They can be addressed by incorporating ML components, such as NNs, directly into the system dynamics of the physics-based model.

With the introduction of NNs, the necessity arises to train the model on the available measurement data F_{meas} . The model parameters θ_{NN} and θ_{ph} are optimised to minimise a defined loss function \mathcal{L} (bottom-middle panel in Fig. 1), which is based on the error between the model output F_{UDE} and the measurement data F_{meas} , for instance, defined by the mean squared error (MSE). The loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{MSE}}(y_{\text{UDE}}(u^{(i:i+l)}, t), y_{\text{meas}}^{(i:i+l)}) \quad (5)$$

where the model is evaluated over an input sequence $u^{(i:i+l)}$, starting at time step i and spanning a window of length l . To obtain y_{UDE} , the system dynamics f are integrated using an ODE solver, which numerically solves the differential equation from the initial state $x(t_i)$ over the time interval $[t_i, t_i + l]$:

$$x([i, i+l]) = \text{ODESolve}(u([i, i+l]), x(t_i), f, [t_i, t_i+l], \theta_{ph}, \theta_{NN}), \quad (6)$$

$$y_{\text{UDE}}([i, i+l]) = g(u([i, i+l]), x([i, i+l]), [t_i, t_i+l], \theta_{ph}, \theta_{NN}). \quad (7)$$

The result of this process is the UDE model (last text-box in Fig. 1) that approximates the measured behaviour, Fig. 1(B). This model is defined by both the physics-based relations and its ML components. It is trained on the measurement data to reduce the error between the UDE model output and the actual measurement data.

The following discussion explains how UDEs satisfy the traits A1–A3, outlined in Section 1. Realisation of A1: Starting with a state-of-the-art mechanistic suspension model ensures that known, basic dynamics are included. By embedding NNs directly into the state-space representation, the model is able to learn and capture additional nonlinearities that influence the system dynamics. Realisation of A2: Literature on UDE design suggests different architectural choices. To ensure interpretability of the learning results, NNs are focused here that introduce nonlinearities into physically interpretable parameters of the mechanistic model, instead of learning deterministic state residuals. Realisation of A3: In UDE learning, mini-batching is employed to enable computationally efficient training on large datasets. This approach allows training

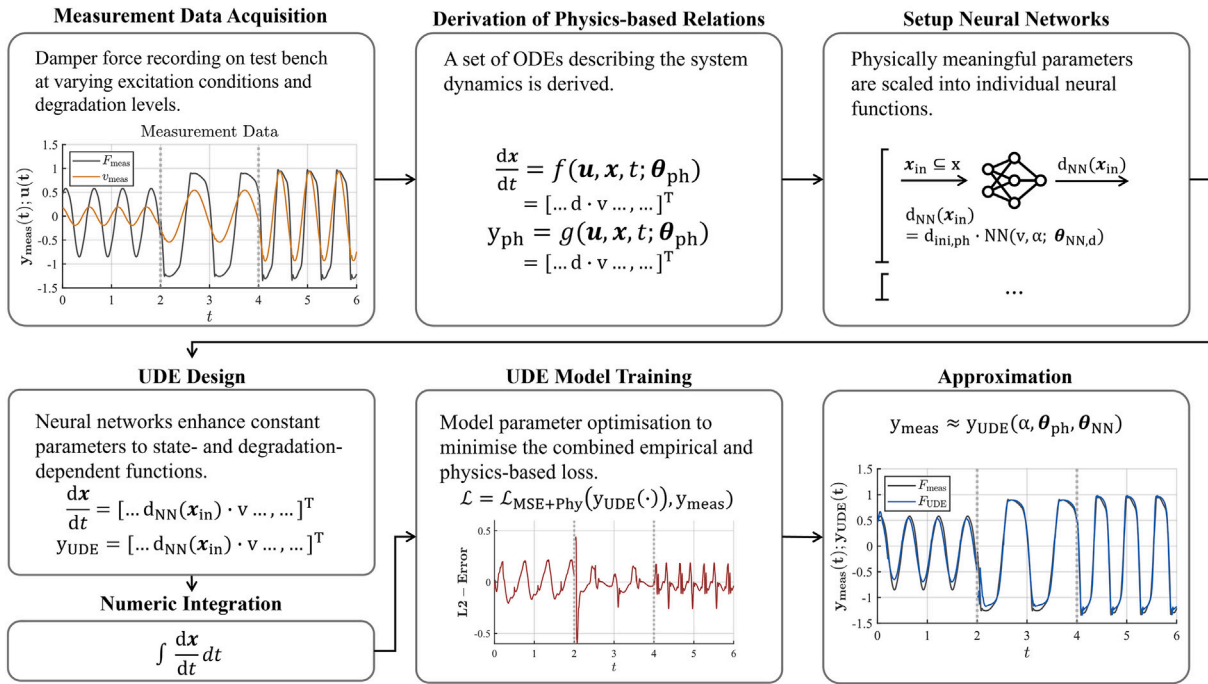


Fig. 1. Overview of the proposed UDE modelling workflow for suspension dampers. Starting from test bench measurements, a physics-based mechanistic model is derived from first principles, whose physically meaningful parameters are enhanced by integrating neural networks as state- and degradation-dependent multiplicative scalars. The resulting UDE model is trained by minimising a combined empirical and physics-based loss $\mathcal{L}_{\text{MSE+Phy}}$. The bottom right panels show the measured damping force (black) and the UDE model estimate (blue) after training, alongside the corresponding estimation error over time.

on multiple datasets with different operating and system conditions. The differential-programming framework PyTorch [29] and a simple third- and fourth-order Runge–Kutta solver are used to numerically solve the UDE [30].

To learn a UDE model based on an appropriate mechanistic model, extensive data were collected using a damper test bench. This data is analysed to support the choice of the mechanistic model.

3. Test bench measurements and analysis

3.1. Experimental setup

The damper under consideration is a passive twin tube hydraulic damper with an internal valve, as shown in Fig. 2.

The damper test bench used to conduct the measurement runs is illustrated in Fig. 3. The cylinder of the damper is firmly screwed to the displacement-exciting piston of the hydropulse system at the bottom via a bracket. The damper piston rod is connected to a ball joint to prevent unnecessary bending moments on the damper. The load cell at the top of the frame measures the force exerted by the damper. Input displacement




Fig. 2. Simplified schematic of a twin-tube shock damper.




Fig. 3. Damper test bench setup.

and velocity are measured by additional sensors at the bottom of the setup.

In total, three distinct measurement runs were conducted using a standard automotive damper test bench [15]. During these experiments, the stroke amplitude was varied between 0.05 m and 0.01 m. For each measurement run, the test bench was controlled to reach certain target velocities at the zero crossing of the excitation amplitude: 0.052 m/s, 0.131 m/s, 0.262 m/s, 0.393 m/s, 0.524 m/s, and 1.048 m/s, which led to variations in the corresponding excitation frequencies. Damper force and stroke amplitude were measured for each experiment with a 10 kHz sampling frequency. Damper velocity and acceleration were calculated through numerical differentiation. To mimic damper degradation, the damper oil was reduced in two steps: one step with a 15% reduction and

Table 1
Test bench measurement configuration.

Parameter	Value / Description
Degradation levels	0%, 15%, 30% oil reduction
Stroke amplitudes	0.05 m, 0.03 m, 0.01 m
Target velocities	0.052, 0.131, 0.262, 0.393, 0.524, 1.048 m/s

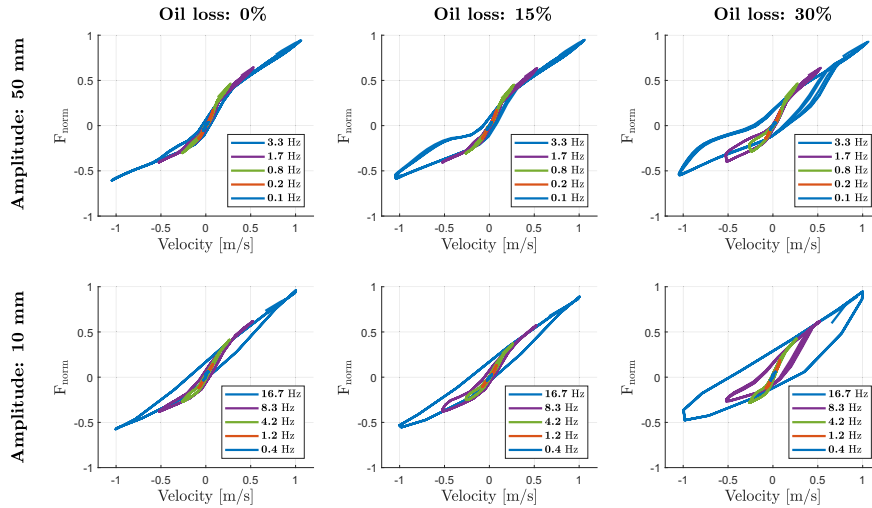


Fig. 4. (Top) normalised force–velocity loops measured with well-functioning and degraded dampers for oil-loss levels of 0%, 15%, and 30% at excitation amplitudes of 50 mm (top row) and 10 mm (bottom row).

another with a 30% reduction. A complete overview of the measurement configuration is provided in Table 1.

Note that the stroke amplitude was adjusted to smaller values, in comparison to the common value of 0.025 m [15] to excite the hysteresis behaviour of the damper.

3.2. Measurement results

Fig. 4 presents normalised force–velocity loops measured on the damper test bench under three oil-loss conditions (0%, 15%, and 30% by volume) for two excitation amplitudes: 50 mm and 10 mm. Each loop is plotted at a series of loading frequencies: 0.1, 0.2, 0.8, 1.7, and 3.3 Hz for 50 mm amplitude, and 0.4, 1.2, 4.2, 8.3, and 16.7 Hz for 10 mm amplitude.

In the absence of oil loss (0%), the hysteresis loops are relatively narrow and nearly symmetric for both amplitudes, with only slight widening as frequency increases. Introducing 15% oil depletion causes the loops to broaden and become modestly skewed: the force during deceleration deviates from that during acceleration, indicating an increase in nonlinear damping. When oil loss reaches 30%, the loops widen substantially, especially at higher frequencies, and exhibit pronounced asymmetry between upstroke and downstroke. This behaviour demonstrates that the damper’s force output depends not only on the instantaneous velocity but also on the velocity history, and that energy dissipation per cycle grows both with excitation frequency and oil depletion. Moreover, reducing the amplitude from 50 mm to 10 mm forces the damper to reach the same peak velocities at higher frequencies; consequently, the 10 mm loops at mid- to high frequencies (e.g., 4.2 Hz and above) are markedly more open than their 50 mm counterparts, revealing that smaller-stroke, higher-frequency conditions exacerbate nonlinear hysteresis effects, particularly when oil loss is significant.

In addition, Fig. 5 shows the difference between the baseline (0% oil loss) normalised force, $F_{norm,0\%}$, and the degraded force, $F_{norm,loss}$, for oil-loss levels of 15% and 30%, again at amplitudes of 10 mm and 50 mm.

At 15% oil loss and 10 mm amplitude, deviations are minimal at low frequency (0.4 Hz) but grow increasingly pronounced at higher frequencies, with the loops shifting upward or downward depending on the direction of motion. At 30% oil loss and 10 mm amplitude, the difference loops at 8.3 Hz and 16.7 Hz reach their largest magnitudes, underscoring that small-stroke, high-frequency operation is most sensitive to degradation. For 50 mm amplitude, 15% oil loss produces modest deviations that become visible above 0.8 Hz; at 30% oil loss, the difference loops at

1.7 Hz and 3.3 Hz display the largest area, indicating that large-stroke, high-frequency excitations suffer substantial force reduction when oil is lost. Across all amplitudes and oil-loss levels, higher frequencies amplify the force deviation. Furthermore, the asymmetric shape of the difference loops reveals a nonlinear, direction-dependent loss of damping capability.

In summary, both figures illustrate that the damper’s hysteresis widens and becomes more asymmetric with increasing oil loss; that smaller amplitudes (and hence higher frequencies for a given peak velocity) magnify nonlinear damping effects; and that degradation has an increasing effect on damping-forces as excitation frequency and amplitude rise. The results show that a damper model that is able to take degradation into account needs to map nonlinear, transient, and asymmetric hysteresis behaviour.

4. Equivalent mechanical model

The scope of the damper model is to include nonlinearities and dynamics within a frequency range of up to 30 Hz. As shown in Fig. 2, a twin-tube damper is designed with several chambers and valves. While the oil flows from one chamber to another through the valve restrictions, a damping-force is generated.

To derive the EMM, this study builds on established approaches by discretising the damper model into three submodels: the pressure-force model, the flow-force model, and the friction-force model. The pressure-force model estimates a force resulting from the damper chamber pressures; the flow-force model estimates a force resulting from the fluid flow through the valve; and the friction-force model estimates a force resulting from the mechanical friction. These models are combined in series or parallel based on Newton’s third law of motion into a full model.

4.1. Equivalent mechanical model for the flow-force

The flow model relates the total oil flow Q_{tot} between the chambers as a function of the pressure drop Δp_{tot} . To derive the model equations, first principles from physics are utilised; in this case, Bernoulli’s equation is applied. Therefore, it is assumed (for a quasi-static damper velocity v_{st}) that the total volume flow is related to the damper velocity v_{st} and A_{eff} , the effective cross-sectional area of the damper by

$$Q_{tot} = v_{st} A_{eff}. \quad (8)$$

The pressure drop is further modelled as a static function of the total volume flow. For details, see Appendix A. The corresponding flow-force

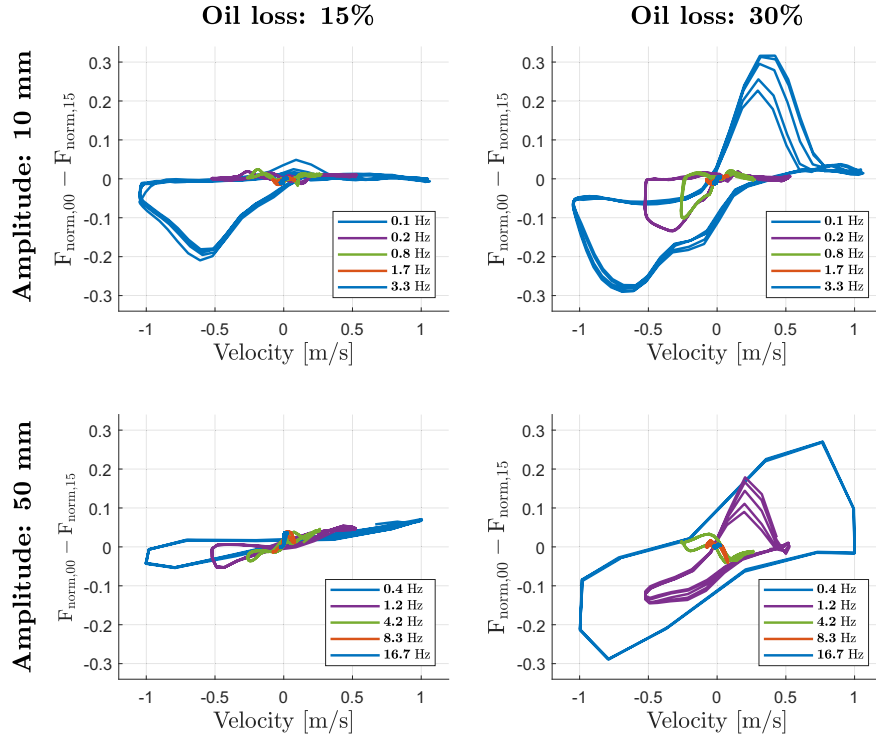


Fig. 5. Difference loops $F_{\text{norm},0\%} - F_{\text{norm},\text{loss}}$ for oil-loss levels of 15% (left column) and 30% (right column) at amplitudes of 10 mm (top row) and 50 mm (bottom row).

F_{fl} is then obtained from the resulting difference between the rebound and compression chamber pressures, Δp_{tot} :

$$F_{\text{fl}} = \Delta p_{\text{tot}} A_{\text{eff}} \quad (9)$$

Combining Eqs. (8) and (9), the flow-force can be rewritten as a function of the quasi-static damper velocity v_{st}

$$F_{\text{fl}} = f_{\text{fl}}(v_{\text{st}}; \theta_{\text{fl}}) \quad (10)$$

with the parameter set θ_{fl} containing parameters that describe the damper geometry, flow characteristics, and fluid properties. The equivalent mechanical model is then described as

$$b_{\text{fl}} = f_{\text{fl}}(\cdot)/v_{\text{st}}, \quad (11)$$

with the damping function b_{fl} .

Focusing on oil loss as a degradation mechanism, let $\alpha \in [0, 1]$ be the normalised level of degradation from 0 (no loss) to 1 (100% oil loss), it is assumed that degradation does not change the geometric and mechanical properties, and thus the relation between the damper velocity v_{st} and the volume flow Q_{tot} . However, degradation can influence and change the non-geometric characteristics, such that θ_{fl} changes to θ_{fl}^* . For instance, due to oil loss, a mixing of gas and oil can occur, which changes the fluid characteristics. Therefore, (9) results in

$$b_{\text{fl}} = f_{\text{fl}}(v_{\text{st}}; \alpha; \theta_{\text{fl}}^*)/v_{\text{st}}. \quad (12)$$

4.2. Equivalent mechanical model for the pressure-force

The pressure-force model describes the dynamic behaviour of the damper through a set of equations that are physically related to the adiabatic compression of the gas present in the damper, the compressibility of the oil, and the compliance of the cylinder. The mechanical equivalent is

$$\begin{aligned} F_{\text{pr}} &= f_{\text{pr},k}(x_{\text{pr}}; \theta_{\text{pr},k}) + f_{\text{pr},b}(\dot{x}_{\text{pr}}; \theta_{\text{pr},b}), \\ k_{\text{pr}} &= f_{\text{pr},k}(\cdot)/x_{\text{pr}}, \quad b_{\text{pr}} = f_{\text{pr},b}(\cdot)/\dot{x}_{\text{pr}}. \end{aligned} \quad (13)$$

with x_{pr} , the relative displacement of the damper and its time derivative \dot{x}_{pr} and the stiffness and damping-force model, referred to as $f_{\text{pr},k}$ and g_b , respectively. The corresponding functions of damping and stiffness are denoted b_{pr} and k_{pr} . For details see Appendix B.

Considering the influence of the degradation α and focusing on oil loss, it is assumed that degradation does not change the geometric properties. Moreover, as mentioned above, mixing of gas and oil in the chambers can occur, leading to oil foam. These effects can influence the overall stiffness and damping characteristics of the model. Therefore, the mechanical equivalents result in

$$\begin{aligned} b_{\text{pr}} &= f_{\text{pr},k}(x_{\text{pr}}; \alpha; \theta_{\text{pr},k}^*)/\dot{x}_{\text{pr}}, \\ k_{\text{pr}} &= f_{\text{pr},b}(\dot{x}_{\text{pr}}; \alpha; \theta_{\text{pr},b}^*)/x_{\text{pr}}. \end{aligned} \quad (14)$$

Remarks on the equivalent mechanical submodels

The presented force models incorporate the following simplifications: It is assumed that the damper is heated to a constant working temperature and thus temperature effects are omitted (viscosity and density depend on the temperature in the chambers). Hydrodynamic friction is neglected and the effect of compliance in the damper mount is not implemented in the model. Moreover, the stiffness and the damping of the submodels are assumed to be constant.

Therefore, the challenge is to determine the functions f_{fl} , $f_{\text{pr},k}$ and $f_{\text{pr},b}$ and extend the model with NNs without violating the underlying physics-based relations. Since oil loss is the primary degradation effect, it is assumed that damping described by the flow-force model decreases with degradation. For the pressure-force model, it is assumed that damping and stiffness decrease with degradation as well.

4.3. Combined equivalent mechanical model

For the combination of the flow-force model and the pressure-force model, the functions that describe stiffness and damping are considered. The pressure submodels k_{pr} and b_{pr} act in parallel and are connected in series with the flow damping b_{fl} . To account for the time delay of the




Fig. 6. Schema of the combined EMM.

valves, a first-order lag element (with Laplace operator s) is used to approximate the lag between \dot{x}_{pr} and v_{st} with time constant T_{pr} ,

$$v_{st} = \frac{1}{T_{pr}s + 1}(\dot{x}_{da} - \dot{x}_{pr}). \quad (15)$$

Balancing the forces and summarising the corresponding equations lead to the differential equations for the nonlinear damper model with input $\mathbf{u} = [\dot{x}_{da}]$ and damper state $\mathbf{x} = [x_{pr}, v_{st}]^T$ with relative displacement x_{pr} and the quasi-static damper velocity v_{st} at the main valve,

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} -\frac{k_{pr}}{b_{pr}+b_{fl}}x_{pr} + \frac{b_{fl}}{b_{pr}+b_{fl}}\dot{x}_{da} \\ -\frac{1}{T_{pr}}v_{st} + \frac{k_{pr}}{T_{pr}(b_{pr}+b_{fl})}x_{pr} + \frac{b_{pr}}{T_{pr}(b_{pr}+b_{fl})}\dot{x}_{da} \end{bmatrix}. \quad (16)$$

The system parameters are $\theta_{phy} = [\theta_{pr,k}, \theta_{pr,b}, \theta_{fl}, T_{pr}]$.

With the solution $\mathbf{x} = [x_{pr}, v_{st}]^T$ of the differential equations at time t the damping-force F_{phy}

$$F_{phy} = \frac{k_{pr} b_{fl}}{b_{pr} + b_{fl}} x_{pr} + \frac{b_{pr} b_{fl}}{b_{pr} + b_{fl}} \dot{x}_{da} + r(x_{da}) \quad (17)$$

is derived with an additional friction submodel r . A schema of the EMM is shown in Fig. 6. Since this paper deals with a twin-tube passive damper, the friction force is low and therefore the friction submodel is neglected here.

Remarks on limitations of the equivalent mechanical model

The EMM can be regarded as a simplified version of a physics-based model, considering the basic physical mechanisms. While the simplifications enable faster computation and easier parameterisation, some essential physical effects may not be captured. Moreover, the EMM lacks effects from degradation, which may not be known and remain unmodelled. Unmodelled physical effects and model simplifications limit the generalisation capability of the EMM across a wide range of damper states and operating conditions.

5. Neural equivalent mechanical model

The enhancement of the EMM with the UDE method to a Neural Equivalent Mechanical Model (NEMM) will now be addressed.

5.1. NEMM model structure

To improve model performance and accuracy, Feed Forward NNs are integrated into the previously shown EMM, which then becomes a UDE model. Specifically, the pressure- and flow-force submodels described in Sections 4.1 and 14 are enhanced using a novel form of residual learning. This involves adding data-driven functions to capture the residual forces in each submodel. These functions are designed as scaling factors, dependent on parts of the EMM state, that are related to the specific model element. The functions are implemented using NNs, hereafter referred to as neural functions. For example, let an arbitrary mechanical stiffness be denoted by k and its corresponding deflection by x . The stiffness is assumed to depend on degradation α and deflection x and therefore described by

$$k(\alpha, x) = k_{ini} \cdot \underbrace{NN_{sc,k}(\alpha, x; \theta_{sc,k})}_{\text{Residual Neural Scaler}} \quad (18)$$

where $NN_{sc,k}(\cdot)$ is a state-dependent function, modelled by a NN, scaling the initial scalar stiffness value k_{ini} . The advantage of this approach is that physical interpretability is preserved, as $k(\alpha, x)$ still retains its physical meaning as a stiffness parameter with $k > 0$. This approach differs from existing methods in the literature, e.g., [27], which add a residual additive term to the system dynamics. The formulation as a residual scaler is beneficial when it comes to EMM parameters of different scales. In the mechanical model of the damper, the value of the pressure-model stiffness is 100 times larger than the damping coefficient of the flow path. Since NNs train more stably for outputs on normalised scales, the approach of additive residuals is limited in performance.

For the application of the UDE methodology, a two-step procedure is followed: First, the EMM model is formulated with constant parameters $(b_{fl,ini}, k_{pr,ini}, b_{pr,ini})$, which are, in the second step, turned into neural functions $(b_{NN,fl}(\cdot), k_{NN,pr}(\cdot), b_{NN,pr}(\cdot))$. These functions correspond to the damping of the flow-force submodel b_{fl} the stiffness k_{pr} and the damping of the pressure-force submodel b_{pr} ,

$$\begin{aligned} b_{NN,fl}(\cdot) &= b_{fl,ini} \cdot NN_{sc,fl}(v_{st}, \alpha; \theta_{sc,fl}), \\ b_{NN,pr}(\cdot) &= b_{pr,ini} \cdot NN_{sc,pr,b}(\alpha; \theta_{sc,pr}), \\ k_{NN,pr}(\cdot) &= k_{pr,ini} \cdot NN_{sc,pr,k}(x_{pr}, \alpha; \theta_{sc,pr}). \end{aligned} \quad (19)$$

The neural networks are embedded into the differential equations derived in Section 4.3 as follows,

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} -\frac{k_{NN,pr}}{b_{NN,pr}+b_{NN,fl}}x_{pr} + \frac{b_{NN,fl}}{b_{NN,pr}+b_{NN,fl}}\dot{x}_{da} \\ -\frac{1}{T_{pr}}v_{st} + \frac{k_{NN,pr}}{T_{pr}(b_{NN,pr}+b_{NN,fl})}x_{pr} + \frac{b_{NN,fl}}{T_{pr}(b_{NN,pr}+b_{NN,fl})}\dot{x}_{da} \end{bmatrix} \quad (20)$$

and

$$F_{phy} = \frac{k_{NN,pr} b_{NN,fl}}{b_{NN,pr} + b_{NN,fl}} x_{pr} + \frac{b_{NN,pr} b_{NN,fl}}{b_{NN,pr} + b_{NN,fl}} \dot{x}_{da} + r(x_{da}). \quad (21)$$

For learning, the NEMM is treated as a black-box model with the additional input of the utilised NNs, which modifies the formal description of the EMM to

$$\frac{d\mathbf{x}}{dt} = f(\mathbf{u}, \mathbf{x}, t, \text{NNs}; \theta), \quad (22)$$

$$F_{NEMM} = g(\mathbf{u}, \mathbf{x}, t, \text{NNs}; \theta),$$

with θ denoting the joint parameter set of physical and NNs-related parameters.

5.2. Proposed NEMM learning

For learning the NEMM, three different types of losses are considered. First, an empirical loss that measures the discrepancy between the model

estimations and the ground truth data, and second, a regularisation loss. The MSE is chosen for the loss function and the L2 norm of the network weights W as the regularisation term,

$$\mathcal{L} = \frac{1}{n} \sum_{i=1}^n (F_{\text{NEMM},i} - F_{\text{meas},i})^2, \quad (23)$$

$$\lambda R(W) = \lambda \|W\|_2,$$

where $\{(F_{\text{NEMM},i}, F_{\text{meas},i})\}_{i=1}^n$ is the set of training instances F_{meas} and corresponding estimations F_{NEMM} .

To incorporate the physics-based knowledge of the EMM into training, key physical relationships, which are abstracted from the derivation of the EMM, are employed. The two components of this physical prior are presented in the next section. These enforce physically consistent estimations and together constitute the third type of loss. Using this physics-based loss along with the empirical loss and regularisation terms in the learning objective, the complete joint learning approach is

$$\begin{aligned} \arg \min_{\text{NN}_{\text{sc},0-n}} & \underbrace{\mathcal{L}(F_{\text{NEMM}}, F_{\text{meas}})}_{\text{Empirical Loss}} + \sum_{i=0}^n \lambda \underbrace{R(\text{NN}_{\text{sc},i})}_{\text{Regularisation Loss}} \\ & + \sum_{i=0}^n \lambda_{\text{phy}} \underbrace{\mathcal{L}_{\text{phy}}(F_{\text{NEMM}})}_{\text{Physical Inconsistency Loss}} \end{aligned} \quad (24)$$

with n utilised neural networks $\text{NN}_{\text{sc},n}$.

5.2.1. Physical inconsistency loss for b_{fl}

From the EMM it is evident that the damping-force monotonically increases with damping velocity v_{st} and satisfies at two different velocities, $v_{\text{st},1}$ and $v_{\text{st},2}$, for the same degradation level α

$$F_{\text{fl}}[v_{\text{st},1}, \alpha] - F_{\text{fl}}[v_{\text{st},2}, \alpha] \leq 0 \quad \text{if } v_{\text{st},1} < v_{\text{st},2}. \quad (25)$$

As no parameter in the above physics-based equations increases with degradation, the damping F_{fl} related to the flow-force model must monotonically decrease with α

$$\frac{F_{\text{fl}}[\alpha_1, v_{\text{st}}]}{v_{\text{st}}} - \frac{F_{\text{fl}}[\alpha_2, v_{\text{st}}]}{v_{\text{st}}} \geq 0 \quad \text{if } \alpha_1 < \alpha_2. \quad (26)$$

On any pair of consecutive velocity values, $v_{\text{st},i+1}$ and $v_{\text{st},i}$ ($v_{\text{st},i+1} > v_{\text{st},i}$), and consecutive degradation values α_{i+1} and α_i ($\alpha_{i+1} > \alpha_i$), the difference in the force estimates can be calculated as

$$\Delta_{\text{vel}}[v_{\text{st}}, \alpha] = \hat{F}_{\text{fl}}[v_{\text{st},i}, \alpha] - \hat{F}_{\text{fl}}[v_{\text{st},i+1}, \alpha], \quad (27)$$

$$\Delta_{\alpha}[v_{\text{st}}, \alpha] = \frac{\hat{F}_{\text{st}}[\alpha_i, v_{\text{st}}]}{[v_{\text{t}}]} - \frac{\hat{F}_{\text{st}}[\alpha_{i+1}, v_{\text{st}}]}{[v_{\text{st}}]}. \quad (28)$$

A positive Δ_{vel} or a negative Δ_{α} constitutes a violation in physical understanding. Let $f_{\text{NN},\text{fl}}$ denote the neural function that estimates the flow-force \hat{F}_{fl} and

$$\hat{F}_{\text{fl}} = f_{\text{NN},\text{fl}}(\cdot) = b_{\text{NN},\text{fl}}(\cdot) v_{\text{st}}. \quad (29)$$

Thus, for differentiable F_{fl} , the global monotonicity conditions

$$\frac{\partial f_{\text{NN},\text{fl}}(\alpha, v_{\text{st}})}{\partial v_{\text{st}}} > 0, \quad (30)$$

$$\frac{\partial (f_{\text{NN},\text{fl}}(\alpha, v_{\text{st}})/v_{\text{st}})}{\partial \alpha} \leq 0y \quad (31)$$

have to be fulfilled. In practice,

$$\mathcal{X}_{\text{st}} = \{\mathbf{u}_{\text{fl}}^{(i,j)} : i = 1, \dots, n_{v_{\text{st}}}, j = 1, \dots, n_{\alpha}\}$$

is considered a uniform grid of $n_{v_{\text{st}}}$ velocity values and n_{α} degradation levels and the two global monotonicity measures

$$U_v := \min_{\mathbf{u}_{\text{fl}} \in \mathcal{X}_{\text{st}}} \frac{\partial f_{\text{NN},\text{fl}}(\mathbf{u}_{\text{fl}})}{\partial v_{\text{st}}}, \quad (32)$$

$$U_{\alpha} := \min_{\mathbf{u}_{\text{NN},\text{fl}} \in \mathcal{X}_{\text{st}}} \frac{\partial (f_{\text{NN},\text{fl}}(\mathbf{u}_{\text{fl}})/v_{\text{st}})}{\partial \alpha} \quad (33)$$

are introduced.

If $U_v \geq 0$ and $U_{\alpha} \leq 0$, then the neural flow-force model f_{fl} satisfies the required global monotonicity. Sampling \mathbf{u}_{fl} over \mathcal{X}_{st} and penalizing any violation yields the physics-based loss

$$\begin{aligned} R_{\text{ph},\text{fl}}(f_{\text{NN},\text{fl}}) = \mathbb{E}_{\mathbf{u}_{\text{fl}} \sim \mathcal{X}_{\text{st}}} & \left[\max\{0, -\partial_{v_{\text{st}}} f_{\text{NN},\text{fl}}(\mathbf{u}_{\text{fl}})\}^2 \right. \\ & \left. + \max\{0, \partial_{\alpha} (f_{\text{NN},\text{fl}}(\mathbf{u}_{\text{fl}})/v_{\text{st}})\}^2 \right]. \end{aligned} \quad (34)$$

5.2.2. Physical inconsistency loss for $b_{\text{NN},\text{pr}}$ and $k_{\text{NN},\text{pr}}$

From the EMM it is evident that the forces related to damping increase monotonically with velocity and forces related to the stiffness increase monotonically with displacement, while both damping and stiffness decrease as degradation increases.

For the pressure-force model, let $f_{\text{NN},\text{pr},k}$ and $f_{\text{NN},\text{pr},b}$ denote the neural functions that estimate the stiffness-force $\hat{F}_{\text{pr},k}$ and the damping-force $\hat{F}_{\text{pr},b}$ related to the pressure-force model, respectively. Then:

$$\hat{F}_{\text{pr},b} = f_{\text{NN},\text{pr},b}(\cdot) = b_{\text{NN},\text{pr}}(\cdot) \dot{x}_{\text{pr}}, \quad (35)$$

$$\hat{F}_{\text{pr},k} = f_{\text{NN},\text{pr},k}(\cdot) = k_{\text{NN},\text{pr}}(\cdot) x_{\text{pr}}. \quad (36)$$

For differentiable functions, the global monotonicity conditions have to be fulfilled for the stiffness-force model

$$\frac{\partial f_{\text{NN},\text{pr},k}(\cdot)}{\partial x_{\text{pr}}} > 0, \quad (37)$$

$$\frac{\partial (f_{\text{NN},\text{pr},k}(\cdot)/x_{\text{pr}})}{\partial \alpha} \leq 0,$$

and for the damping-force model

$$\frac{\partial f_{\text{NN},\text{pr},b}(\cdot)}{\partial v_{\text{pr}}} > 0, \quad (38)$$

$$\frac{\partial (f_{\text{NN},\text{pr},b}(\cdot)/\dot{x}_{\text{pr}})}{\partial \alpha} \leq 0.$$

This results in the physics-based loss for the stiffness-force model

$$\begin{aligned} R_{\text{ph},\text{pr},k}(f_{\text{NN},\text{pr},k}) = \mathbb{E}_{\mathbf{u}_{\text{pr}} \sim \mathcal{X}_{\text{pr}}} & \left[\max\{0, -\partial_{x_{\text{pr}}} f_{\text{NN},\text{pr},k}(\cdot)\}^2 \right. \\ & \left. + \max\{0, \partial_{\alpha} (f_{\text{NN},\text{pr},k}(\cdot)/x_{\text{pr}})\}^2 \right], \end{aligned} \quad (39)$$

and for the damping-force model

$$\begin{aligned} R_{\text{ph},\text{pr},b}(f_{\text{NN},\text{pr},b}) = \mathbb{E}_{\mathbf{u}_{\text{pr}} \sim \mathcal{X}_{\text{pr}}} & \left[\max\{0, -\partial_{\dot{x}_{\text{pr}}} f_{\text{NN},\text{pr},b}(\cdot)\}^2 \right. \\ & \left. + \max\{0, \partial_{\alpha} (f_{\text{NN},\text{pr},b}(\cdot)/\dot{x}_{\text{pr}})\}^2 \right]. \end{aligned} \quad (40)$$

5.3. Training and test data

For learning the neural functions, the test-bench setup remains the same as the data set described in Section 3. However, for the flow-force model, only specifically selected training samples are considered. This selection process ensures that only quasi-static damper states are included. Table 2 shows an overview of the used data.

Table 2
Train, validation, and test data configuration.

Configuration	Train & Val. Data	Test Data
Stroke amplitude	10 & 50 mm	10 & 50 mm
Excitation frequencies	0.1 - 16.7 Hz	0.1 - 16.7 Hz
Oil loss	0 & 30%	15%

Table 3
Hyperparameter Search Space for Optimisation.

Hyperparameter	Options
Batch size	{4, 8, 16}
Hidden size	{8, 16, 32}
Number of layers	{3, 5, 7, 9}
Learning rate	{0.5e-2, 1e-2, 0.1e-2}

Flow-force model: For measurements of the quasi-static damper characteristics, the damper force is measured along with the corresponding damper displacement and velocity. The excitation is performed at different frequencies. To determine the quasi-static characteristic, the individual velocity values are used as inputs. However, the excitation with sinusoidal strokes causes hysteresis, which is then removed for the sake of identification. Therefore, the maximum damper force within a range of $\pm 5\%$ of the displacement amplitude around $x = 0$ m is used; the maximum velocity occurs at $x = 0$ m.

Pressure-force model: For training, validation, and testing of the pressure-force model, the three datasets described earlier are used, and further reduced to decrease the training effort.

From the measurements $X_i = [v_{st,i}, \alpha_i]$ and the corresponding single output $y_i = F_{meas,i}$, the data were extracted and used to train the model.

To ensure that the input data is represented on a comparable scale, a feature-wise normalisation was applied. Specifically, absolute maximum scaling was used, where each feature value X_i was divided by the maximum absolute value of that feature in the training dataset. Due to the varying lengths and generally long time series of the gathered measurement data, the data was cut into shorter sequences. This ensured efficient use of the available data and prevented excessively long training times caused by long time series. These sequences had a duration of 0.5 s, long enough to capture frequency effects above 2 Hz.

5.4. Implementation details and hyperparameter tuning

The neural functions ($b_{NN,\beta}(\cdot)$, $k_{NN,pr}(\cdot)$, $b_{NN,pr}(\cdot)$) estimate a continuous scalar. Each NN is implemented as an MLP with L hidden layers of H neurons, where each layer consists of a linear transformation, LayerNorm, LeakyReLU, and dropout. The final linear layer is initialised so that its bias yields an output of 0.5 as the first output. The models are trained using the Adam optimiser [31]. For the L2 loss, λ is set to 0.001.

To further optimise performance, the structural hyperparameters were tuned. Table 3 summarises the search space for each hyperparameter:

Each candidate configuration was evaluated by observing the validation RMSE. The optimal setup consisted of a batch size of 16, a hidden size of 16, five layers, and a learning rate of 0.5e-2.

6. Model evaluation and knowledge discovery

To investigate the effect of different architectural choices on model quality, four model variants are considered. The variants are selected to provide a structured, ablation-style evaluation of the Neural-EMM's individual contributions, progressing from a standard industry baseline to the proposed NEMM architecture:

(1) **Static Empirical Model (SEM):** This model relates the damper velocity to the damper force through a nonlinear force-velocity mapping. The map is determined by measuring damper forces at quasi-static velocity states on a test bench, which serve as support points. This

procedure is standard practice in automotive damper modelling; see [15] for details.

(2) **Linear EMM:** This variant employs the ODE formulation of the EMM to describe the dynamic coupling between damper force and velocity. All model parameters are treated as constants, i.e., without any state dependency. It represents the simplest physics-based dynamic baseline, isolating the benefits of introducing nonlinear modelling capacity through neural network integration.

(3) **NEMM-Light:** Building on the Linear EMM model, this variant introduces a single NN to model the damping of the flow-force submodel b_{fl} as a neural function of quasi-static velocity v_{st} and degradation level α . Compared to the full NEMM, this variant isolates the contribution of the neural pressure-force submodel parameterisation of damping b_{pr} and stiffness k_{pr} .

(4) **NEMM:** The full NEMM extends NEMM-Light by incorporating two additional NNs to model the damping b_{pr} and the stiffness k_{pr} of the pressure-force submodel as neural functions of internal model states damper velocity \dot{x}_{da} and the degradation level α . It isolates the contribution of a reduced neural network architecture within the UDE framework, directly quantifying the effect of network capacity on performance relative to the full NEMM.

6.1. Model evaluation and comparison

To compare the four model architectures w.r.t. their model quality, a quantitative analysis is conducted: Each model is fitted or trained using the training dataset, and its generalisation performance is subsequently evaluated on the test dataset. Table 4 presents the root-mean-squared error (RMSE) of each model on the test dataset.

The SEM (considered as baseline) achieves an RMSE of 0.077. In comparison, the linear EMM formulation worsens performance (RMSE = 0.120, 155.8% relative error), whereas introducing a single NN in NEMM-Light reduces the error to RMSE = 0.057 (74.0% relative error), and the full NEMM yields the best performance (RMSE = 0.050, 64.9% relative error).

Since the NEMM formulation demonstrates the best overall performance, it is examined in greater detail through a qualitative analysis. For this purpose, the measured and modelled hysteresis loops are compared at two excitation frequencies and two oil-loss levels. Fig. 7 presents the results: the normalised measured forces $F_{meas, norm}$ (top row) and the corresponding NEMM force-estimations F_{NEMM} (bottom row) are plotted against the normalised damper velocity \dot{x}_{da} for excitation frequencies of $f = 8.3$ Hz (blue curves) and $f = 16.7$ Hz (red curves), both at an amplitude of $A = 10$ mm. The left column corresponds to 0% oil loss, whereas the right column represents 30% oil loss.

As oil loss increases from 0% to 30%, both the measured and estimated loops exhibit increased hysteresis and nonlinearity. In all four plots, the NEMM estimations (bottom row) qualitatively match the measured hysteresis loops (top row). At 0% oil loss and 8.3 Hz, the nearly straight measured loop is reproduced almost exactly by the UDE model. At 0% oil loss and 16.7 Hz, the slightly wider measured loop is matched well, both slope and area, by the UDE model. When oil loss reaches 30%, the measured loops at 8.3 Hz and 16.7 Hz become significantly wider and more nonlinear; in each case, the NEMM captures the increased loop area, nonlinear progression, and the small loading-unloading asymmetry at 16.7 Hz. For all analysed scenarios,

Table 4

Comparison of the performance of the four model architectures on the test dataset.

Model	Model error (RMSE)	Rel. Model error (%)
SEM	0.077	100.0
Linear EMM	0.120	155.8
NEMM-Light	0.057	74.0
NEMM	0.050	64.9

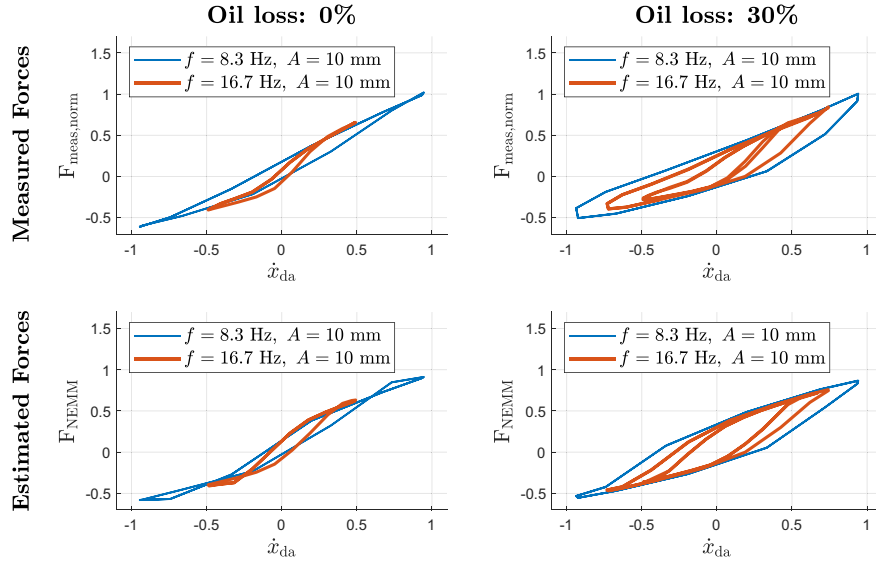


Fig. 7. Comparison of measured and NEMM-estimated hysteresis loops at two excitation frequencies ($f = 8.3$ Hz in blue and $f = 16.7$ Hz in red, both at amplitude $A = 10$ mm). The top row shows the normalised measured forces $F_{\text{meas, norm}}$ versus normalised velocity \dot{x}_{da} [m/s] for (left) 0% oil loss and (right) 30% oil loss.

a discrepancy remains between the modelled and measured asymmetry of the hysteresis loops with respect to the damper velocity \dot{x}_{da} : the measured asymmetry is larger than the modelled one.

6.2. Evaluation of physical inconsistency loss

To evaluate the influence of the physical inconsistency loss, introduced in Section 5.2, the NEMM is trained both with and without the consistency terms, and the resulting model quality is assessed. Furthermore, all learned neural functions ($b_{\text{NN,fl}}(\cdot)$, $k_{\text{NN,pr}}(\cdot)$, $b_{\text{NN,pr}}(\cdot)$) for the damping of the flow–force submodel (b_{fl}) the stiffness (k_{pr}) and the damping of the pressure-force submodel (b_{pr}), are examined in detail.

First, the flow–force submodel is examined. For this purpose, the weighting factor of the physical inconsistency loss, λ_{ph} , is set to 0 and 0.025, an empirically determined value. The joint training loss is obtained by combining (34) and (24) as

$$\mathcal{L} = \text{MSE}(\hat{F}_{\text{fl}}, F_{\text{meas, static}}) + \lambda R(b_{\text{NN,fl}}(\cdot)) + \lambda_{\text{ph}} R_{\text{ph,fl}}(b_{\text{NN,fl}}(\cdot)), \quad (41)$$

with \hat{F}_{fl} estimated according to (29).

To assess and illustrate the learned mapping, the model was stimulated with synthetic inputs covering the input domain ($v_{\text{st}} \in [-1.2, 1.2]$ and $\alpha \in [0, 1]$). Fig. 8 shows the results: the estimated nonlinear damping

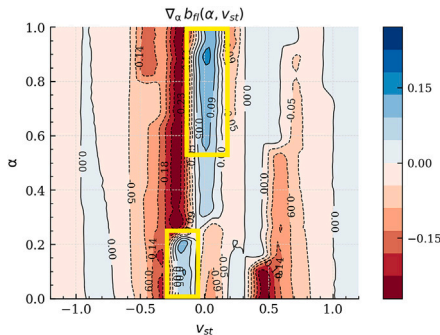


Fig. 8. Exclusion of the physical inconsistency loss term with $\lambda_{\text{ph}} = 0.0$: Illustration of b_{fl} as a function of degradation α and quasi-static velocity v_{st} [m/s], modelled by the learned neural function $b_{\text{NN,fl}}(\cdot)$. Violations of physics-based assumptions are marked in yellow.

b_{fl} as a function of the quasi-static velocity (v_{st}) and degradation (α) for $\lambda_{\text{ph}} = 0$. Omitting the physical inconsistency loss term causes violations (highlighted in yellow) of the physical relations. Fig. 9 shows the result for $\lambda_{\text{ph}} = 0.025$, where the physical inconsistency loss enforces $\nabla_{\alpha} b_{\text{NN,fl}} \geq 0$, eliminating such violations.

Next, the neural functions $b_{\text{NN,pr}}(\cdot)$, and $k_{\text{NN,pr}}(\cdot)$, related to the pressure-force model, are examined. For the learning loss formulation, (40) and (39) are combined with (24):

$$\mathcal{L} = \text{MSE}(\hat{F}_{\text{NEMM}}, F_{\text{meas}}) + \lambda [R(k_{\text{NN,pr}}(\cdot)) + R(b_{\text{NN,pr}}(\cdot))] + \lambda_{\text{ph}} [R_{\text{ph,b}}(k_{\text{NN,pr}}(\cdot)) + R_{\text{ph,k}}(b_{\text{NN,pr}}(\cdot))], \quad (42)$$

with \hat{F}_{NEMM} estimated according to (22).

For both neural functions it is observed that omitting the physical inconsistency loss penalty leads to violations of the monotonicity requirements, while its inclusion eliminates such violations.

To investigate the influence of the physical inconsistency loss on the NEMM, a parameter study is conducted by systematically varying the weighting factor λ_{ph} . The resulting test performances are summarised in Table 5. These results show that incorporating the physical inconsistency loss improves the overall generalisation performance of the full NEMM model.




Fig. 9. Inclusion of the physical inconsistency loss term with $\lambda_{\text{ph}} = 0.025$: Illustration of b_{fl} as a function of degradation α and quasi-static velocity v_{st} [m/s] modelled by $b_{\text{NN,fl}}(\cdot)$. No violations of physics-based assumptions.

Table 5
Evaluation of the NEMM model performance for varying values of the physical inconsistency loss weighting factor λ_{ph} .

λ_{ph}	Model error (RMSE)	Rel. Model error (%)
0	0.069	100.0
0.005	0.049	71.0
0.015	0.051	73.9
0.025	0.052	75.4

6.3. Effects of degradation on neural equivalent mechanical model

Owing to the model’s ability to accurately represent the degraded damper behaviour, the learned dependency on degradation is analysed in more detail. Therefore, each learned neural function is stimulated with sampled data, spanning its respective input domain.

Starting with the neural function related to the damping of the flow-force model (b_{fl}) the sampled data spans the degradation range $\alpha \in [0, 1]$ and quasi-static damper velocity $v_{st} \in [-1, 1]$. The results are shown as a contour plot in Fig. 10, top right subplot B. The plot reveals that b_{fl} captures a nonlinear coupling between degradation and velocity. Furthermore, it can be observed that damping decreases with an increasing degradation level α . The top left subplot A illustrates this relationship for the specific velocity of -0.5 m/s. Here, the damping drops from approximately 0.8 to 0.75, corresponding to a reduction of about 6.25%. Velocity has a more pronounced effect on damping. This is clearly shown in the two lower subplots: the upper one C displays damping versus velocity, and the lower one D shows the corresponding force. The results indicate that damping is particularly high at low velocities.

Next, the neural functions $k_{NN,pr}(\cdot)$ and $b_{NN,pr}(\cdot)$ related to the pressure-force model are investigated. To stimulate these two components, synthetic data spanning the degradation range $\alpha \in [0, 1]$ is used for b_{pr} . For k_{pr} , the displacement range $x_{pr} \in [-0.003, 0.003]$ is additionally used. The results are illustrated in Figs. 11 and 12, for b_{pr} and k_{pr} , respectively. Both b_{pr} and k_{pr} decrease with increasing degradation. The damping coefficient b_{pr} decreases from approximately 3 to 1, corresponding to a reduction of about 66.7%. The stiffness k_{pr} decreases from around 400 to 200, which corresponds to a 50% reduction. The symmetry of k_{pr} w.r.t. x_{pr} is not affected by degradation, indicating that

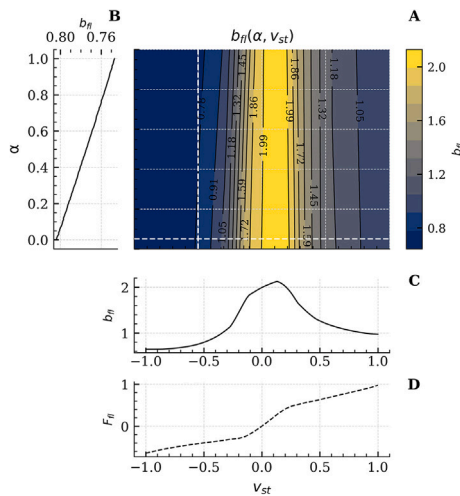


Fig. 10. Damping of the flow-force submodel: Illustration of the learned mapping of b_{fl} [Ns/m] as a function of degradation α and quasi-static velocity v_{st} [m/s], subplot A, modelled by the neural function $b_{NN,fl}(\cdot)$. Subplot B illustrates b_{fl} at $v_{st} = -0.5$ as a function of the degradation level, while subplot C shows b_{fl} at $\alpha = 0.9$ as a function of the quasi-static velocity. The bottom subplot D presents the corresponding damping force $F_{fl}(v_{st})$.




Fig. 11. Damping of the pressure-force submodel: Illustration of the learned mapping of b_{pr} [Ns/m] as a function of the degradation level α , modelled by the neural function $b_{NN,pr}(\cdot)$.

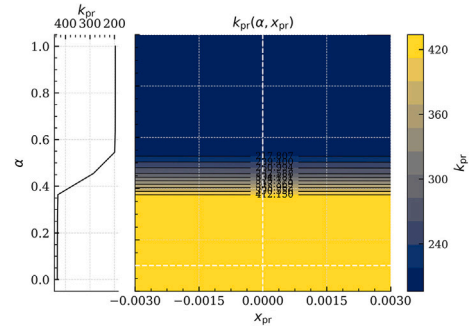


Fig. 12. Stiffness of the pressure-force submodel: Illustration of the learned mapping of k_{pr} [N/m] as a function of displacement x_{pr} [m] and degradation level α , modelled by the neural function $k_{NN,pr}(\cdot)$.

the stiffness is reduced equally in both compression and rebound directions. Both the reduction in stiffness and damping of the pressure-force model are physically plausible, as lower degradation levels correspond to more foamed oil.

7. Summary and conclusions

Degraded vehicle suspension dampers can significantly impact vehicle safety and ride comfort. Simulation models of dampers, that handle degradation effects, can improve the understanding of changing characteristics and allow for effect analysis on the component and vehicle levels. To analyse the effect of degradation, various real-world measurements of a hydraulic suspension damper were performed. These measurements indicate that increasing degradation introduces strongly nonlinear, transient, and asymmetric changes in the damper dynamics.

Equivalent Mechanical Models (EMMs), typically used for developing component models in vehicle dynamics simulations, generalise well across various operating conditions. However, the analysis conducted has shown that the simplified design of an appropriate EMM model for well-functioning dampers leads to discrepancies compared to empirical observations when degraded dampers are considered.

An alternative approach to EMMs is the application of Universal Differential Equations (UDEs). This modelling approach facilitates the integration of Neural Networks into EMMs, increasing the mappability of nonlinear and transient relationships and turning them into hybrid models. The method was used to design a hybrid model, termed Neural EMM, which is capable of modelling dynamics, including those arising from degradation.

The Neural EMM, compared to the EMM, exhibits enhanced accuracy, evidenced by reduced estimation discrepancies across a wide range of damper states and degradation conditions. The model was evaluated across excitation frequencies from 0.1 to 16.7 Hz, stroke amplitudes of 10 mm and 50 mm, and a degradation level of 15% oil loss as an unseen test condition, demonstrating the model’s ability to generalise to unknown degradation states.

The proposed framework is not limited to a specific type of damper or degradation effect. The EMM structure incorporates fundamental physical principles that apply more broadly to hydraulic dampers. For example, the flow-force submodel relates oil flow between chambers to

pressure drop using Bernoulli's equation, without assuming particular internal geometry, valve layout or fluid rheology. Consequently, in principle, the same model topology can describe a broad class of hydraulic dampers beyond the twin-tube configuration studied here. Similarly, the UDE approach is independent of the specific parameterisation of the residual dynamics: the embedded neural networks act as universal approximators for any smooth behaviour not captured by the physics-based submodel, whether arising from changes in seal friction, viscosity variation or other degradation-induced phenomena. Physics-based auxiliary losses also ensure that the learned functions remain physically consistent, regardless of the specific degradation effect present.

Although the UDE modelling approach showed promising results by combining the strengths of physics-based and data-driven models, the ability of the trained model to generalise is limited by the diversity of the training data. Operating conditions that were not represented during training, such as extreme temperature ranges, highly transient excitation profiles or severe degradation states that go beyond what was observed, may reduce modelling accuracy. Moreover, the ability of the model to map a variety of degradation effects, such as seal wear, elastomeric ageing and valve fretting fatigue, to the damper dynamics is not focused on here. A systematic evaluation of robustness across a broader range of damper configurations, degradation effects and operating ranges is identified as a task for future work.

CRedit authorship contribution statement

Lorenz Ott: Writing – review & editing, Writing – original draft, Software, Methodology, Data curation, Conceptualization. **Silke Redecker:** Writing – review & editing, Visualization. **Torben Gräber:** Supervision, Conceptualization. **Michael Unterreiner:** Supervision. **Johannes Edlmann:** Writing – review & editing, Supervision. **Manfred Plöchl:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Programme.

Appendix A. Derivation of the flow-force model

For the valve flow Q_{tot} in (8), three elements are considered in the rebound phase: leak restriction, blow-off valve, port restriction; and one element in the compression phase: intake valve, [13]. For the pressure drop

$$\Delta p = B Q^r, \quad r \in \mathbb{R} \quad (\text{A.1})$$

as a function of the flow rate Q , Reybrouck [32] suggested to model B as a function of a pressure loss factor K and viscosity ν :

$$\Delta p_{leak} = B_{leak} Q_{leak}^{1.75} = K_{leak} \nu^{0.25} Q_{leak}^{1.75}, \quad (\text{A.2})$$

$$\Delta p_{port} = B_{port} Q_{port}^{1.75} = K_{port} \nu^{0.25} Q_{port}^{1.75}. \quad (\text{A.3})$$

The blow-off valve includes a tension spring and a flat disc. The valve remains closed until the blow-off pressure is larger than the pre-loaded spring pressure ($\Delta p_{blow-off} \geq \Delta p_0$). Therefore, the blow-off valve is modelled as a function of the spring stiffness K_{spring} , the opening height y of the valve, and the cross-sectional area A_v . There holds,

$$K_{spring} y = \Delta p_{blow-off} A_v - \Delta p_0 A_v. \quad (\text{A.4})$$

Neglecting inertial, viscous and momentum terms, the valve dynamics simplify to the static equation for the blow-off valve, [12]:

$$\bar{K}_{spring} Q_{blow-off} = (\Delta p_{blow-off} - \Delta p_0) \sqrt{\Delta p_{blow-off}}. \quad (\text{A.5})$$

For the total valve characteristics, the sub-elements are connected either in series or in parallel. Therefore, the pressure drop across two elements in series is the sum of the pressure drops across the individual elements, and the pressure drop across two elements in parallel is the same as the drops across the individual elements. Further, the total flow of two valves in parallel is equal to the sum of the flows through the individual elements, see [11].

$$\Delta p_{parallel} = \Delta p_{blow-off} = \Delta p_{leak}, \quad (\text{A.6})$$

$$\Delta p_{total} = \Delta p_{port} + \Delta p_{parallel}, \quad (\text{A.7})$$

$$Q_{total} = Q_{port} = Q_{leak} + Q_{blow-off}. \quad (\text{A.8})$$

Finally, the total pressure flow Q_{total} sums up to

$$\Delta p_{total} = K_{port} \nu^{0.25} Q_{total}^{1.75} + K_{leak} \nu^{0.25} Q_{leak}^{1.75}, \quad (\text{A.9})$$

$$Q_{tot} = Q_{leak} + \quad (\text{A.10})$$

$$\frac{(K_{leak} \nu^{0.25} Q_{leak}^{1.75} - \Delta p_0)(K_{leak} \nu^{0.25} Q_{total}^{1.75})^{0.5}}{\bar{K}_{spring}}. \quad (\text{A.11})$$

with K_{port} the stiffness of the port restriction, K_{leak} the stiffness of the leak restriction, ν the fluid viscosity, Q_{total} the total flow rate, Q_{leak} the flow rate through the leak restriction, Δp_{total} the resulting total pressure drop, and Δp_0 the preload pressure drop.

Appendix B. Derivation of the pressure-force model

The compressibility of the gas and therefore the pressure in the reserve tube is approximated from a polytropic compression and polytropic expansion during the rebound stroke. As the damper is moving, it is assumed that no heat is transferred from the volume of gas, and the gas force changes adiabatically, [11],

$$p_{gas,dyn} V_{gas,dyn}^\gamma = p_{gas,static} V_{gas,static}^\gamma. \quad (\text{B.1})$$

γ , the polytropic exponent, is set to 1.4, a common assumption for high-frequency changes. Finally,

$$p_{pr} = p_{pr,0} \left(\frac{V_{pr,gas,0}}{V_{pr,gas,0} + A_{rod} x_{pr}} \right)^\gamma, \quad (\text{B.2})$$

and the pressure force in (13) results in

$$F_{pr} = F_{pr,0} \left(\frac{V_{pr,gas,0}}{V_{pr,gas,0} + A_{rod} x_{pr}} \right)^\gamma. \quad (\text{B.3})$$

with p_{pr} the gas pressure $p_{pr,0}$ the initial pressure in the reserve tube, $V_{pr,gas,0}$ the initial gas volume, A_{rod} the damper-rod area, x_{pr} the piston displacement, and F_{pr} , $F_{pr,0}$ the resulting pressure force and its initial value, respectively.

Data availability

The data that has been used is confidential.

References

- [1] Unger A, Pellegrini E, Henning K, Lohmann B. A model for dynamic feedforward control of a semi-active damper. In: Proc. of the 5th international conference on integrated modelling and analysis in applied control and automation (IMAACA 2011). vol. 164, 2011.
- [2] Mastinu G, Ploechl M, Road and offroad vehicle system dynamics handbook, CRC Press, Boca Raton, 2014, <https://doi.org/10.1201/b15560>
- [3] TÜV Süd, Die häufigsten mängel. 2024. <https://www.tuvsud.com/de-de/publikationen/tuev-report> [Accessed: 16 January 2024].
- [4] Zwosta T, Kubenz J, Prokop G. Experimental analysis of the behavior of automotive twin-tube dampers degraded by loss of oil and pressure. SAE Technical Paper 2023-01-5084, SAE International, 2023, <https://doi.org/10.4271/2023-01-5084>

- [5] Bedük MD, Çalışkan K, Henze R, Küçükay F, Effects of damper failure on vehicle stability, *Proc Inst Mech Eng D J Automob Eng* 2013;227(7):1024–39, Publisher: Sage Publications Sage UK: London, England. <https://doi.org/10.1177/0954407012465242>
- [6] Vaculín O, Svoboda J, Valášek M, Steinbauer P, Influence of deteriorated suspension components on ABS braking, *Veh Syst Dyn* 2008;46(S1):969–79.
- [7] Hu S, Yang M, Meng D, Hu R. Damping performance of the degraded fluid viscous damper due to oil leakage. In: *Structures*. vol. 48, Elsevier; 2023. p. 1609–19.
- [8] Schramm T, Zwosta T, Prokop G, Investigation and phenomenological modeling of degraded twin-tube shock absorbers for oil and gas loss, *Vehicles* 2025;7(1).
- [9] Basargan H, Mihály A, Gáspár P, Sename O. Fault-tolerant semi-active suspension control for degradation in damper performance. In: *29th Mediterranean conference on control and automation (MED)*. 2021. <https://doi.org/10.1109/MED51440.2021.9480288>
- [10] Dai L, Chi M, Guo Z, Gao H, Wu X, Sun J, Liang S, A physical model-neural network coupled modelling methodology of the hydraulic damper for railway vehicles, *Veh Syst Dyn* 2023;61(2):616–37.
- [11] Duym SW, Simulation tools, modelling and identification, for an automotive shock absorber in the context of vehicle dynamics, *Veh Syst Dyn* 2000;33(4):261–85. [https://doi.org/10.1076/0042-3114\(200004\)33:41-UFT261](https://doi.org/10.1076/0042-3114(200004)33:41-UFT261)
- [12] Segel L, Lang H, The mechanics of automotive hydraulic dampers at high stroking frequencies, in: *The dynamics of vehicles on roads*, Routledge, 2018, pp. 194–214.
- [13] Duym SW, Stiens R, Baron GV, Reybrouck KG. Physical modeling of the hysteretic behaviour of automotive shock absorbers. Tech. rep., SAE Technical Paper, 1997.
- [14] Lee K, Numerical modelling for the hydraulic performance prediction of automotive monotube dampers, *Veh Syst Dyn* 1997;28(1):25–39. <https://doi.org/10.1080/00423119708969347>
- [15] Reimpell J, Stoll H, Betzler J, *The automotive chassis: engineering principles*, Elsevier, 2001.
- [16] Fash J. Modeling of shock absorber behavior using artificial neural networks. Tech. rep., SAE Technical Paper, 1994.
- [17] Rackauckas C, Ma Y, Martensen J, Warner C, Zubov K, Supekar R, Skinner D, Ramadhan A, Edelman A, Universal differential equations for scientific machine learning, Aug, arXiv preprint arXiv:2001.04385, 2020. <https://arxiv.org/abs/2001.04385>.
- [18] Raissi M, Perdikaris P, Karniadakis GE, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J Comput Phys* 2019;378:686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [19] Li Q-Q, Xu Z-D, Zhao Q, Dong Y-R, Wang X-W, Liu T-F, Physics-guided surrogate modeling of dynamic mechanical properties of viscoelastic dampers considering micromolecular chain structures, *Structures* 2025;74:108533, <https://doi.org/10.1016/j.istruc.2025.108533>
- [20] Li Q-Q, Xu Z-D, Dong Y-R, He J-X, Tian Y, He Z-H, Guo Y-Q, Mechanical property tests and physics-informed data-driven modeling of viscoelastic materials subjected to thermal-oxidative aging, *Constr Build Mater* 2024;414:134920, <https://doi.org/10.1016/j.conbuildmat.2024.134920>
- [21] Zeng T, Wang T, Zeng Z, Zhang F, Byeon J, Wang Y, Zou Y, Wang Y, Jiao J, Claudel C, Chen X, Damper-B-PINN: damper characteristics-based Bayesian physics-informed neural network for vehicle state estimation, arXiv preprint arXiv:2502.20772, 2025; <https://arxiv.org/abs/2502.20772>.
- [22] Kamp T, Ultsch J, Brembeck J. Closing the sim-to-real gap with physics-enhanced neural ODEs. In: *Proceedings of the 20th international conference on informatics in control, automation and robotics*. Rome, Italy: SCITEPRESS - Science and Technology Publications; 2023. p. 77–84, <https://doi.org/10.5220/0012160100003543>
- [23] Delgado-Trujillo S, Alvarez DA, Bedoya-Ruiz D, Hysteresis modeling of structural systems using physics-guided universal ordinary differential equations, *Comput Struct* 2023;280:106988, <https://doi.org/10.1016/j.compstruc.2023.106988>
- [24] Ajay A, Wu J, Fazeli N, Bauza M, Kaelbling LP, Tenenbaum JB, Rodriguez A. Augmenting physical simulators with stochastic neural networks: case study of planar pushing and bouncing. In: *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE; 2018. p. 3066–73, <https://doi.org/10.1109/IROS.2018.8593995>
- [25] Viana FA, Nascimento RG, Dourado A, Yucesan YA, Estimating model inadequacy in ordinary differential equations with physics-informed neural networks, *Comput Struct* 2021;245:106458, <https://doi.org/10.1016/j.compstruc.2020.106458>
- [26] Yucesan YA, Viana FA, Manin L, Mahfoud J, Adjusting a torsional vibration damper model with physics-informed neural networks, *Mech Syst Signal Process* 2021;154:107552, <https://doi.org/10.1016/j.ymsp.2020.107552>
- [27] Kuzhiyil JA, Damoulas T, Widanage WD, Neural equivalent circuit models: universal differential equations for battery modelling, *Appl Energy* 2024;371:123692, <https://doi.org/10.1016/j.apenergy.2024.123692>
- [28] ElGazzar A, Gerven MV, Universal differential equations as a common modeling language for neuroscience, Mar, arXiv preprint arXiv:2403.14510, 2024.
- [29] Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, others, PyTorch: an imperative style, high-performance deep learning library, *Adv Neural Inf Process Syst* 2019;32.
- [30] Chen RT, Rubanova Y, Bettencourt J, Duvenaud DK, Neural ordinary differential equations, *Adv Neural Inf Process Syst* 2018;31.
- [31] Kingma DP, Mohamed S, Jimenez Rezende D, Welling M, Semi-supervised learning with deep generative models, *Adv Neural Inf Process Syst* 2014;27.
- [32] Reybrouck K, A non linear parametric model of an automotive shock absorber, *SAE Trans* 1994; 1170–7. <https://doi.org/10.4271/940869>