Proceedings of 17th International Symposium on Advanced Vehicle Control (AVEC'26)
7-11 September 2026, Tsukuba, Japan

A Port-Hamiltonian-Inspired Identification of Longitudinal Vehicle Dynamics

Alireza AGHASHARIF[1)], Gioele DEFRANCESCO[1)], Gastone Pietro ROSATI PAPINI[1)]

1) Department of Industrial Engineering, University of Trento

Abstract: In this note, we develop a port-Hamiltonian-inspired model-structured neural network (MSNN) for longitudinal vehicle-dynamics identification. In the MSNN paradigm, part of the computational graph is fixed from physical principles, while trainable neural modules are assigned only to unknown governing equations. The proposed model follows this principle and extends an existing structured framework for longitudinal vehicle dynamics by imposing an explicit mechanical energy balance and nonnegative dissipative force channels. This work does not propose a fully free port-Hamiltonian neural state-space model. Instead, it embeds the longitudinal acceleration identification problem inside a fixed mechanical storage and momentum-balance structure. This makes the mechanical power balance hold by construction, rather than through an added loss penalty. The model is trained using measured velocity, acceleration, driveline actuation, braking, gear selection, and altitude, without calibrated wheel-force measurements, drivetrain maps, aerodynamic coefficients, rolling-resistance constants, or controlled dynamometer tests. The learned components are therefore interpreted as effective longitudinal force terms consistent with the available telemetry and the imposed port-Hamiltonian structure, while the model identifies the vehicle's longitudinal acceleration from the measured data.

Keywords: Vehicle Dynamics Theory; Modeling, Identification, and Estimation

1. Introduction

Classical vehicle dynamics are commonly modeled by organizing the dynamics of the vehicle and its interaction with the environment as a finite set of differential equations. This is not only a mathematical formulation, but also an engineering modeling choice, since it defines which physical effects are retained, aggregated, or neglected. Over the last few decades, neural networks have become a mainstream modeling tool based on the universal approximation theory [(1)]. However, purely data-driven neural networks may have limited ability to learn and generalize the long-term behavior of dynamical systems governed by known physical principles. In recent years, this has motivated the emergence of scientific machine learning, where physical knowledge is incorporated into the learning process, for example by penalizing violations of the governing physics [(5,8)].

In vehicle dynamics, this motivates neural models whose structure reflects the physical domain rather than arbitrary input–output maps. For longitudinal dynamics, the structural choice is essential: different force domains naturally impose distinct dissipative or conservative roles. Therefore, the main difficulty is not only to improve prediction accuracy, but also to preserve the physical meaning of the learned force contributions. Recent advances have embedded mechanical properties into the network architecture, such as Hamiltonian neural networks, which encode conservative dynamics through an energy function [(4)]. However, due to the dissipative nature of longitudinal vehicle dynamics, a purely ideal Hamiltonian reconstruction would be insufficient. Port-Hamiltonian theory extends Hamiltonian modeling to systems with dissipation and external ports, representing dynamics through stored energy, interconnection, dissipation, and power-conjugate ports [(2,3)]. This makes it a suitable framework for physics-guided neural modeling, where the learned dynamics are restricted to respect the separation between storage, supply, and dissipation. Recent port-Hamiltonian and pseudo-Hamiltonian neural networks have shown how this structure can separate Hamiltonian, damping, and external force contributions in forced and dissipative systems [(7,9)]. A port-Hamiltonian partition into power ports preserves these roles by construction, rather than letting a neural map identify them globally. In this paper, the longitudinal identification problem is formulated from an energy and port-based viewpoint. The vehicle is described through a fixed mechanical storage, whose rate of change is governed by power exchanged through longitudinal force ports. This formulation separates conservative storage effects from externally supplied and dissipated power, so that acceleration prediction is not produced by an unconstrained input–output map, but by a learned force balance embedded in a prescribed mechanical energy structure.

The resulting model can be viewed as a port-Hamiltonian-inspired MSNN: the Hamiltonian core defines the kinetic–potential energy storage and the force–velocity power pair, while the trainable modules identify the unknown power-contributing force channels which enables the direct energy-based interpretation of the power conjugate ports and the dissipativity diagnostics.

The contribution is a grey-box identification framework in which neural learning is restricted to the non-conservative force laws, while the overall model remains constrained by a mechanical power balance. The scope of the study is limited by the available measurements. The identified force channels should therefore be understood as mechanically constrained longitudinal force contributions inferred from the available vehicle-level measurements. Their role is not to reproduce separately calibrated subsystem parameters, but to provide a physically structured decomposition of the measured acceleration through the imposed momentum balance.

2. Problem setting

A finite-dimensional input–state–output port-Hamiltonian system as defined in [(3)] is written as

$$\dot{x} = \left[J(x) - R(x)\right]\nabla H(x) + G(x)u, \tag{1}$$

$$y = G(x)^\top \nabla H(x), \tag{2}$$

where $x \in \mathbb{R}^n$ is the n-dimensional vector of the state, $u \in \mathbb{R}^m$ is the input, $y \in \mathbb{R}^m$ is the output, and $H : \mathbb{R}^n \to \mathbb{R}$ is the scalar Hamiltonian, interpreted as the stored energy. The matrix $J(x) = -J(x)^\top \in \mathbb{R}^{n\times n}$ describes power-conserving interconnection, $R(x) = R(x)^\top \succeq 0 \in \mathbb{R}^{n\times n}$ describes dissipation, and $G(x) \in \mathbb{R}^{n\times m}$ defines the external power ports. Along trajectories of eq. (1),

$$\begin{aligned}\dot H(x) &= \nabla H(x)^\top \dot x \\ &= \nabla H(x)^\top \big[J(x) - R(x)\big]\nabla H(x) + \nabla H(x)^\top G(x)u \\ &= -\nabla H(x)^\top R(x)\nabla H(x) + y^\top u. \end{aligned} \tag{3}$$

The port variables u and y are power-conjugates: their product $y^\top u$ is the instantaneous power supplied through the external ports. The skew-symmetry of $J(x)$ eliminates the internal interconnection term, whereas $R(x) = R(x)^\top \succeq 0$ makes $\nabla H(x)^\top R(x)\nabla H(x)$ a nonnegative dissipated-power term. Therefore,

$$\dot H(x) \le y^\top u. \tag{4}$$

If H is bounded from below, this gives the standard passivity statement with respect to the supply rate $(y^\top u)$. In this work, this relation is not used to define a fully general port-Hamiltonian neural state-space model. Instead, it motivates a vehicle-specific realization in which the mechanical storage and force port are fixed, while the unknown non-conservative effects are represented by structured force channels.

2.1. Fixed-storage longitudinal realization

Longitudinal motion is modeled along a one-dimensional curvilinear road coordinate. Let $s(t) \in [0, L]$ denote the vehicle position along a route of length $L > 0$, and let $v(t) \ge 0$ denote the forward longitudinal velocity. With fixed vehicle mass $M > 0$, define the longitudinal momentum

$$p(t) = Mv(t), \tag{5}$$

and the mechanical state

$$x(t) = \begin{bmatrix} s(t) \\ p(t) \end{bmatrix}. \tag{6}$$

These canonical mechanical coordinates are used because longitudinal forces act directly on the momentum balance.

Let $h : [0, L] \to \mathbb{R}$ be a differentiable road-altitude profile, where $h(s)$ denotes the altitude at road coordinate s. The Hamiltonian is fixed as the total mechanical storage,

$$H(s,p) = \underbrace{\frac{p^2}{2M}}_{\text{kinetic energy}} + \underbrace{Mg\,h(s)}_{\text{potential energy}}. \tag{7}$$

Its gradient is

$$\nabla H(s,p) = \begin{bmatrix} Mgh'(s) \\ p/M \end{bmatrix} = \begin{bmatrix} Mgh'(s) \\ v \end{bmatrix}, \tag{8}$$

where $h'(s)$ is the spatial derivative of the altitude profile. Only $h'(s)$ enters the equations of motion; therefore, adding a constant offset to h shifts the Hamiltonian by a constant but does not change the dynamics, the port variables, or the power balance.

For this one-degree-of-freedom mechanical system, the canonical interconnection matrix is

$$J = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}. \tag{9}$$

The conservative Hamiltonian vector field is therefore

$$J\nabla H(s,p) = \begin{bmatrix} v \\ -Mgh'(s) \end{bmatrix}. \tag{10}$$

Thus $Mgh(s)$ is the gravitational potential-energy term, while $-Mgh'(s)$ is the corresponding longitudinal grade force appearing in the momentum equation.

The longitudinal force port is chosen as

$$G = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \tag{11}$$

The associated collocated port output is

$$y = G^\top \nabla H(s,p) = v. \tag{12}$$

This fixes the force–velocity port used in the momentum balance.

2.2. Force ports and momentum balance

In the proposed model, the mechanical storage H, the interconnection matrix J, and the force-port matrix G are fixed by the longitudinal realization above. We do not learn a free port-Hamiltonian state-space model, nor do we identify a general state-dependent dissipation matrix $R(x)$. Instead, the unknown non-conservative effects are represented as learned longitudinal force channels acting through the prescribed momentum port.

Using the force port in eq. (11) and the collocated output in eq. (12), the port input u_θ has the physical meaning of a longitudinal force, while the port output is the vehicle velocity. Hence,

$$yu_\theta = vu_\theta \tag{13}$$

is the mechanical power exchanged through the prescribed force port. The port variables and the supervised identification output play different roles: the port output is $y = v$, whereas the quantity compared with data is the longitudinal acceleration,

$$\hat a_\theta = \frac{\dot p_\theta}{M}. \tag{14}$$

The scalar force input applied through the momentum port is decomposed as

$$u_\theta = F_{\mathrm{drive},\theta} - F_{\mathrm{road},\theta} - F_{\mathrm{brake},\theta} + F_{\mathrm{res},\theta}, \tag{15}$$

where $F_{\mathrm{drive},\theta}$ is the effective driveline force, $F_{\mathrm{road},\theta}$ is the aggregate road-load force, $F_{\mathrm{brake},\theta}$ is the braking force, and $F_{\mathrm{res},\theta}$ is a bounded signed residual correction. The first three channels define the structured mechanical backbone, whereas the residual is introduced only to compensate for remaining modeling errors.

Substituting eq. (15) into the fixed mechanical realization gives

$$\dot x = J\nabla H(x) + Gu_\theta. \tag{16}$$

Using eqs. (10) and (11), the corresponding momentum equation is

$$\dot p_\theta = -Mgh'(s) + F_{\mathrm{drive},\theta} - F_{\mathrm{road},\theta} - F_{\mathrm{brake},\theta} + F_{\mathrm{res},\theta}. \tag{17}$$

By Newton's second law, the acceleration prediction is

$$\hat a_\theta = \frac{1}{M}\left(-Mgh'(s) + F_{\mathrm{drive},\theta} - F_{\mathrm{road},\theta} - F_{\mathrm{brake},\theta} + F_{\mathrm{res},\theta}\right). \tag{18}$$

Thus, the model is trained as an acceleration predictor, but its prediction is generated from a force-level momentum balance.

Remark 1 (Force-channel dissipation). The dissipative effects are written as explicit force channels rather than forced into a purely state-dependent viscous form $-R(x)\nabla H(x)$. In the present one-dimensional mechanical coordinates, a scalar momentum damping term would have the form

$$F_{\mathrm{loss}}(s,p) = d(s,p)\,v, \qquad d(s,p) \ge 0. \tag{19}$$

This representation is unnecessarily restrictive for experimental vehicle telemetry. Road-load may include velocity-independent rolling resistance, braking is modulated by an external brake command, and driveline effects depend on actuation and operating conditions. In particular, a road-load force with nonzero rolling resistance would require

$$d(v) = \frac{F_{\text{road}}(v)}{v}, \tag{20}$$

which is singular as $v \to 0^+$ whenever $F_{\text{road}}(0) > 0$. The force-channel form therefore preserves the port interpretation while allowing each non-conservative force to depend on the measured variables associated with its physical origin. The required condition for dissipativity is nonnegative extracted power at the force port on the forward-driving domain, not the existence of a state-only viscous damping representation.

By fixing the mechanical storage H, the interconnection J, and the force port G, and by representing the unknown non-conservative effects as learned force channels in the momentum equation, we obtain the following mechanical power identity.

Proposition 1 (Mechanical power identity). For the fixed-storage longitudinal realization defined by eqs. (7), (9), and (11), consider the structured backbone obtained from eq. (17) by setting $F_{\text{res},\theta} = 0$:

$$\dot{p} = -Mgh'(s) + F_{\text{drive},\theta} - F_{\text{road},\theta} - F_{\text{brake},\theta}. \tag{21}$$

Along its trajectories, the Hamiltonian eq. (7) satisfies

$$\dot{H} = vF_{\text{drive},\theta} - vF_{\text{road},\theta} - vF_{\text{brake},\theta}. \tag{22}$$

Proof. From eq. (7),

$$\frac{\partial H}{\partial s} = Mgh'(s), \qquad \frac{\partial H}{\partial p} = \frac{p}{M} = v. \tag{23}$$

Therefore,

$$\dot{H} = Mgh'(s)\dot{s} + v\dot{p}. \tag{24}$$

Using $\dot{s} = v$ and substituting eq. (21) gives

$$\begin{aligned}\dot{H} &= Mgh'(s)v + v\left(-Mgh'(s) + F_{\text{drive},\theta} - F_{\text{road},\theta} - F_{\text{brake},\theta}\right)\\ &= vF_{\text{drive},\theta} - vF_{\text{road},\theta} - vF_{\text{brake},\theta}.\end{aligned} \tag{25}$$

The altitude-dependent terms cancel because the road-slope contribution is generated by the conservative Hamiltonian term $Mgh(s)$.

Corollary 1 (Dissipativity and driveline passivity). On the forward-driving domain $v \geq 0$, assume that

$$F_{\text{road},\theta} \geq 0, \qquad F_{\text{brake},\theta} \geq 0. \tag{26}$$

Then the road-load and braking channels cannot increase the mechanical energy, and

$$\dot{H} \leq vF_{\text{drive},\theta}. \tag{27}$$

This follows directly from eq. (22), since $vF_{\text{road},\theta} \geq 0$ and $vF_{\text{brake},\theta} \geq 0$. Thus, the driveline acts as the signed energy-supply channel, while road-load and braking are dissipative by construction. In this model, the inequality is enforced by fixing the Hamiltonian structure and learning the dissipative effects as sign-constrained force channels, rather than by learning a free dissipation matrix $R(x)$.

For the residual-augmented predictor, the exact power identity contains an additional term $vF_{\text{res},\theta}$. Since $F_{\text{res},\theta}$ is bounded but signed, it is treated as an additional learned correction force and is not included in the hard dissipativity certificate. Consequently, the dissipativity claim applies to the structured road-load and braking channels of the backbone, while the residual contribution is reported separately as a bounded modeling correction.

3. Vehicle dynamics estimation

The data are taken from the Lancia Delta on-road benchmark of Da Lio et al. [6]. The experiment consists of an approximately 50 min, 55.7 km extra-urban drive including straights, curves, hills, and intersections. The vehicle was driven under dry-road conditions, with no reported hard braking, wheel slip, traction-control intervention, or stability-control intervention in the training or validation portions.

Following the benchmark protocol, the first 36.0 km of the trip are used for training and the remaining 19.7 km for validation. The measured signals used in this work are

$$\mathcal{D} = \{v_k, T_k, b_k, G_k, h_k, a_k\}_{k=0}^{N-1}, \tag{28}$$

where v_k is velocity, T_k driveline actuation, b_k brake signal, G_k gear, h_k altitude, and a_k longitudinal acceleration. Longitude and latitude are not used. The benchmark data are filtered and downsampled to 20 Hz. The acceleration signal is used only as the supervised output target. The preceding subsections fix the mechanical structure. Following the model-structured neural-network paradigm of Da Lio et al. [6], the measured telemetry is processed through separate force branches and combined through the momentum balance.

The driveline force is represented by a gear-dependent finite-memory model driven by the recent history of the normalized driveline command,

$$F_{\text{drive},\theta,k} = M\left(w_{\text{d}}^{(G_k)}\right)^{\top}\tilde{T}_k^{[25]}, \tag{29}$$

where $\tilde{T}_k^{[25]}$ collects the previous 25 samples of the normalized driveline signal. The vector $w_{\text{d}}^{(G_k)}$ therefore acts as a finite-impulse-response (FIR) kernel, capturing the dynamic influence of recent driveline commands on the generated traction force. Using gear-specific, G_k parameter vectors allows the model to account for the different driveline characteristics associated with each transmission ratio. The driveline branch is unconstrained in sign and may represent either traction or engine-braking effects.

The road-load force is modeled as

$$F_{\text{road},\theta,k} = M\left(c_0 + c_2\tilde{v}_k^2\right), \qquad c_0, c_2 \geq 0, \tag{30}$$

where c_0 captures an effective rolling-resistance contribution and $c_2\tilde{v}_k^2$ captures the dominant velocity-dependent drag contribution.

The braking force is represented by a nonnegative finite-memory model driven by the brake signal,

$$F_{\text{brake},\theta,k} = Mw_{\text{b}}^{\top}\tilde{b}_k^{[25]}, \qquad w_{\text{b}} \geq 0. \tag{31}$$

Remark 2 (Spatial reconstruction). The fixed-storage formulation uses a spatial road-altitude profile $h(s)$, whereas the dataset provides altitude as a time-indexed sequence h_k. The altitude contribution is computed from the reconstructed spatial altitude profile. Since altitude is measured as a time-indexed signal h_k, the road coordinate is first reconstructed from velocity as

$$s_{k+1} = s_k + \Delta t v_k, \qquad s_0 = 0. \tag{32}$$

The pairs (s_k, h_k) are then interpreted as samples of a spatial altitude profile $\hat{h}(s)$. The corresponding grade contribution is written as

$$a_{\text{h},k} = -g\hat{h}'(s_k), \tag{33}$$

where $\hat{h}'(s_k)$ denotes the estimated dimensionless spatial derivative of the reconstructed altitude profile with respect to road distance. Thus, the altitude term is not learned as an independent scalar correction from h_k; it is obtained from the estimated road grade and corresponds to the continuous port-Hamiltonian grade acceleration $-gh'(s_k)$.

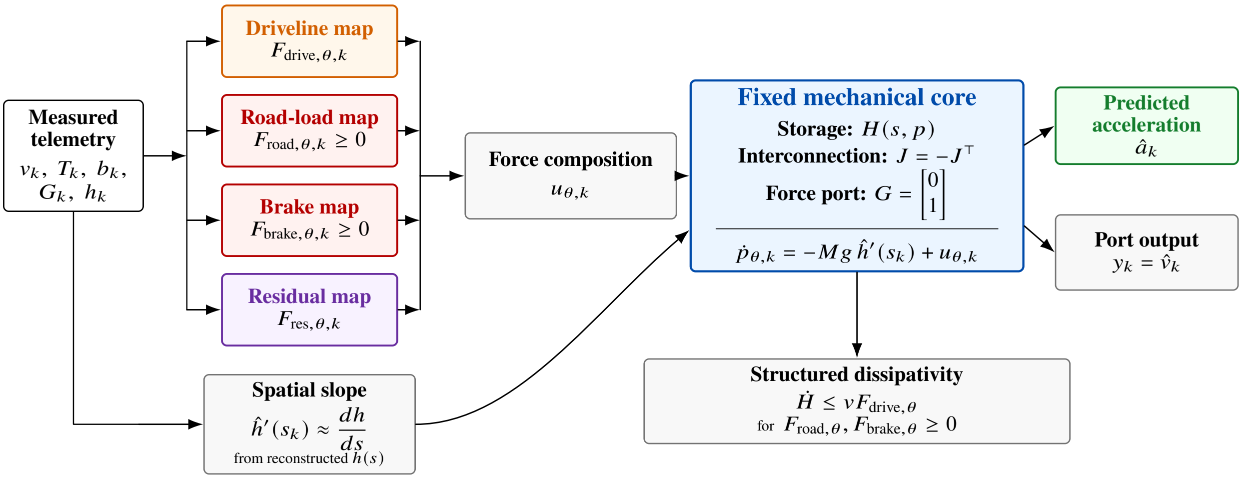

Fig. 1: Structured port-Hamiltonian-inspired vehicle architecture. The architecture is interpreted through the port-Hamiltonian passivity relation $\dot{H} \leq y^\top u$, while the acceleration $\hat{a}_k$ is the supervised prediction output.

3.1. Acceleration target correction

The measured IMU acceleration contains a small constant offset, which causes substantial velocity drift when integrated over the long experiment. Therefore, a scalar bias is estimated using only the official training segment. The corrected target is

$$\bar{a}_k = a_k - b^\star, \tag{34}$$

where

$$b^\star = \arg\min_{b\in\mathbb{R}} \sum_{k\in\mathcal{I}_{\text{tr}}} \left[v_0 + \sum_{j=0}^{k-1} \Delta t\,(a_j - b) - v_k \right]^2 . \tag{35}$$

The same $b^\star$ is applied unchanged to validation. This correction defines only the supervised target and does not introduce acceleration feedback into the model. In the reported experiment, $b^\star = 0.02323\,\text{m/s}^2$. Without this correction, raw acceleration integration over the full record gives a velocity RMSE of $40.33\,\text{m/s}$ and final drift of $69.45\,\text{m/s}$. After correction, the dominant drift is removed. Since full-horizon integration remains sensitive to low-frequency errors, acceleration prediction is kept as the primary benchmark metric. On the $1.25\,\text{s}$ finite-history window used by the model, the local equivalent acceleration inconsistency is $0.0634\,\text{m/s}^2$, slightly below the $0.067\,\text{m/s}^2$ noise RMS reported in [6].

4. Training protocol

The model is identified from the corrected acceleration target $\bar{a}_k$ defined in eq. (34). The acceleration measurement is used only as the supervised output target and is not included among the input features. All normalization constants are computed using the official training segment only.

The training procedure follows the model-structured neural-network (MSNN) principle: the force-based structure of the longitudinal momentum balance is preserved, while trainable maps are used only for the unknown force laws. Training has two stages: first, a structured FIR backbone is fitted; second, a bounded residual MLP is added to compensate for unmodeled effects. The trainable parameters are grouped as

$$\theta = \{\theta_{\text{FIR}}, \theta_{\text{res}}\}, \tag{36}$$

corresponding to the structured backbone and the residual MLP.

In the first stage, let

$$\theta_{\text{FIR}} = \left\{ w_{\text{d}}^{(G)},\ w_{\text{b}},\ c_0,\ c_2,\ \right\} \tag{37}$$

collect the parameters of the gear-dependent driveline FIR kernels, the braking FIR kernel, the road-load coefficients, and the effective local road-geometry branch.

The corresponding structured acceleration prediction is

$$\hat{a}_{\text{FIR},k} = a_{\text{drive},k} + a_{\text{road},k} - a_{\text{brake},k}. \tag{38}$$

Equivalently, since $p = Mv$, this prediction defines the force-level momentum balance

$$\dot{p}_{\text{FIR},k} = M\hat{a}_{\text{FIR},k}. \tag{39}$$

Thus, the model is trained against acceleration data, while the structured prediction remains equivalent to a force-level momentum balance through $\dot{p}_{\text{FIR},k} = M\hat{a}_{\text{FIR},k}$.

The FIR structure makes the backbone linear in θ_{FIR}. On the fitting index set $\mathcal{I}_{\text{fit}}$, the structured prediction can be written in acceleration units as

$$\hat{a}_{\text{FIR}} = \Phi_{\text{FIR}}\theta_{\text{FIR}}, \tag{40}$$

where Φ_{FIR} contains the signed acceleration-level regressors. The road-load columns are $-\tilde{v}_k^2$ and -1, the braking columns are $-\tilde{b}_k^{[n_h]}$, the local road-geometry columns contain η_k and its bias column, and the driveline columns activate only the gear-specific FIR block associated with G_k.

The structured parameters are obtained by solving the bounded regularized least-squares problem

$$\theta_{\text{FIR}}^\star = \arg\min_{\theta_{\text{FIR}}} \frac{1}{|\mathcal{I}_{\text{fit}}|} \sum_{k\in\mathcal{I}_{\text{fit}}} \left(\hat{a}_{\text{FIR},k}(\theta_{\text{FIR}}) - \bar{a}_k\right)^2 + \lambda_{\text{FIR}}\,\|\theta_{\text{FIR}}\|_2^2, \tag{41}$$

subject to

$$c_0 \geq 0, \qquad c_2 \geq 0, \qquad w_{\text{b}} \geq 0. \tag{42}$$

These constraints enforce nonnegative road-load and braking magnitudes, which enter the momentum balance as dissipative force channels. The driveline and road-geometry branches are left signed, allowing traction, engine-braking effects, and uphill/downhill grade-related contributions. In the reported experiment, $\lambda_{\text{FIR}} = 3\times10^{-5}$ is used to improve the conditioning of the structured least-squares

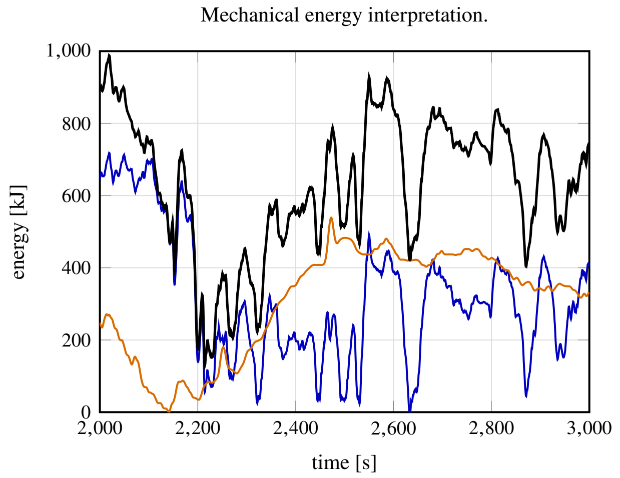

Fig. 2: Mechanical storage decomposition over the validation interval. The blue curve denotes the kinetic energy K, the orange curve denotes the gravitational potential energy U, and the black curve denotes the total mechanical storage $H = K + U$.

fit. In the second stage, a bounded residual correction is added to the structured backbone:

$$\hat{a}_k = \hat{a}_{\mathrm{FIR},k} + \Delta a_{\mathrm{res},k}, \qquad \Delta a_{\mathrm{res},k} = B\tanh\big(f_{\theta_{\mathrm{res}}}(z_k)\big), \tag{43}$$

where z_k is a causal feature vector supplied to the residual MLP, constructed from the measured telemetry histories used by the structured model and excluding measured acceleration. The scalar $B = 0.45\,\mathrm{m/s^2}$ fixes the maximum admissible magnitude of the residual acceleration correction. Equivalently, the residual corresponds to the bounded signed force contribution

$$F_{\mathrm{res},\theta_{\mathrm{res}},k} = M\Delta a_{\mathrm{res},k}. \tag{44}$$

Thus, the residual is not a replacement for the force-based model, but an additional bounded correction force.

The last layer of $f_{\theta_{\mathrm{res}}}$ is initialized to zero, so that initially

$$\Delta a_{\mathrm{res},k} = 0, \qquad \hat{a}_k = \hat{a}_{\mathrm{FIR},k}. \tag{45}$$

During residual training, the structured parameters $\theta^{\star}_{\mathrm{FIR}}$ are kept fixed and only θ_{res} is optimized.

Let s_a denote the standard deviation of the corrected acceleration target $\bar{a}_k$, computed on the official training segment. The residual network is trained by minimizing the normalized acceleration error, together with a residual-magnitude penalty and weight regularization:

$$\begin{aligned}\mathcal{L}_{\mathrm{res}}(\theta_{\mathrm{res}}) = &\frac{1}{|\mathcal{I}_{\mathrm{fit}}|}\sum_{k\in\mathcal{I}_{\mathrm{fit}}}\left(\frac{\hat{a}_{\mathrm{FIR},k} + \Delta a_{\mathrm{res},k} - \bar{a}_k}{s_a}\right)^2 \\ &+ \lambda_\Delta \frac{1}{|\mathcal{I}_{\mathrm{fit}}|}\sum_{k\in\mathcal{I}_{\mathrm{fit}}}\left(\frac{\Delta a_{\mathrm{res},k}}{s_a}\right)^2 + \lambda_{\mathrm{res}}\,\|\theta_{\mathrm{res}}\|_2^2 .\end{aligned} \tag{46}$$

The normalization by s_a makes both acceleration-related terms dimensionless and uses only statistics from the official training segment. The first term fits the corrected acceleration target, the second term discourages unnecessarily large residual corrections relative to the training acceleration scale, and the last term is the residual-network weight regularization.

For model selection, the official training segment is split internally: 0–30 km is used for fitting, while 30–36 km is used only to monitor the residual-network performance and select the stopping

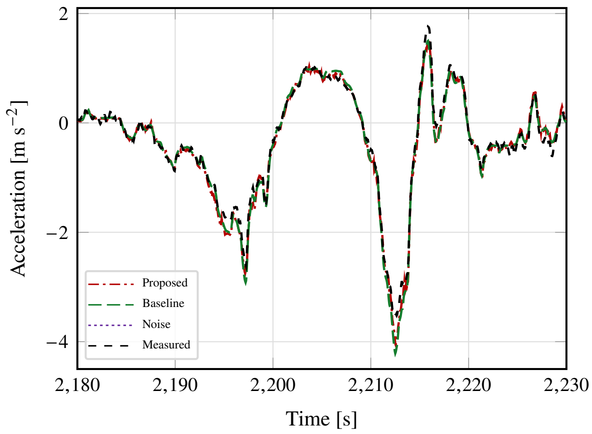

Fig. 3: Example of measured and predicted acceleration over a representative transient interval. The proposed model is shown in red. The noise entry denotes the estimated acceleration-noise level used as a reference in the residual analysis.

epoch. After epoch selection, the structured FIR backbone is refitted on the full official training segment 0–36 km, and a newly initialized residual network is trained on the same segment for the selected number of epochs. The official validation segment is kept unseen during bias estimation, normalization, least-squares fitting, residual training, epoch selection, and final model training.

In the proposed implementation, the finite-memory length is $n_h = 25$, corresponding to 1.25 s at the sampling rate 20 Hz. Therefore, the driveline and braking histories used in the FIR branches are $\tilde{T}_k^{[n_h]}$ and $\tilde{b}_k^{[n_h]}$, respectively. The residual feature vector z_k uses the extended causal setting: driveline history, driveline-increment history, brake history, local altitude features, short velocity history, current velocity, and one-hot gear encoding. The residual MLP has two hidden layers with 64 units per layer and hyperbolic-tangent activations. It is trained with Adam using learning rate 2×10^{-4}, weight-decay coefficient $\lambda_{\mathrm{res}} = 10^{-5}$, mini-batches of 4096 samples, and gradient-norm clipping at 5. The residual regularization weight is $\lambda_\Delta = 5\times10^{-3}$. Internal training is run for at most 180 epochs with patience 80 and random seed 7. The selected epoch is 71, and the final model is trained for this number of epochs before the official validation evaluation.

To support transparency and reproducibility, an open-source implementation of the models and resources are available at https://github.com/tonegas/nnodely-applications/.

5. Results and diagnostics

Figure 3 showcases the validation acceleration comparison of our proposed model, the true acceleration and the legacy result over a selected transient period. Figure 4 demonstrates the Amplitude Spectral Density of the true acceleration, noise, proposed model, and the legacy result. Figure 2 displays the total energy variation of the vehicle in the designed port-based structure over the entire validation horizon. The table 1 summarize the results through the mechanical energy balance.

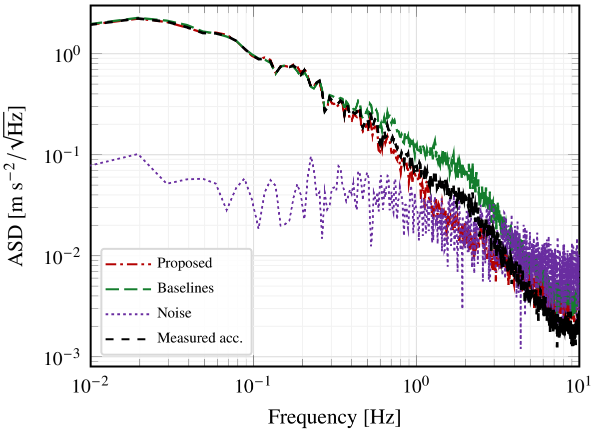

Fig. 4: Amplitude spectral density of the corrected validation acceleration, proposed prediction, established baseline, and estimated noise.

Table 1: Validation acceleration RMSE and port-based diagnostics. m/s^2 (6) benchmark.

Quantity	Value
Validation RMSE	
Da Lio convolutive unstructured (6)	$0.158\,m/s^2$
Da Lio convolutive structured (6)	$0.105\,m/s^2$
Da Lio recurrent unstructured (6)	$0.129\,m/s^2$
Da Lio recurrent structured (6)	$0.115\,m/s^2$
Structured FIR baseline	$0.1087\,m/s^2$
Proposed model	$\mathbf{0.1029}\,m/s^2$
Port-based diagnostics of the proposed model	
Minimum road-load force	81.48 N
Residual power RMS	1.346 kW
Power-identity RMS error	4.89×10^{-4} W

6. Conclusion

A Port-Hamiltonian-Inspired Gray-Box Identification of Longitudinal Vehicle Dynamics was identified from a single on-road dataset using measured vehicle telemetry. The mechanical structure is fixed from first principles, while actuation and dissipation are learned under hard physical constraints. Validation results show accurate longitudinal acceleration prediction across traction and braking phases, the acceleration results are complemented by finite-window consistency checks, while full-horizon velocity reconstruction is treated with caution because it is sensitive to low-frequency acceleration bias. Moreover, the measured mechanical energy variation closely matches the model-implied power balance, confirming consistency with the imposed energy-rate identity. This agreement indicates that the learned force components remain physically coherent and respect nonnegative dissipation when evaluated on experimental data. Future work will extend the proposed framework to coupled longitudinal–lateral vehicle dynamics, including braking transients and tire slip, through an explicit Dirac-structure representation of power-conserving interconnections and dissipation mechanisms of the vehicle.

7. Acknowledgments

We acknowledge financial support under the FIS 2 Call, Grant Assignment Decree No. 1236 adopted on 01/08/2023 by the Italian Ministry of University and Research (MUR), for the project FIS-2023-03684 “Structured neural network framework for modeling and control of autonomous systems – Neu4mes”, CUP E53C24003800001.

8. Use of AI

The authors used Grammarly for language and grammar correction. OpenAI's ChatGPT was used to assist with text editing and polishing.

References

(1) G. Cybenko, Approximation by Superpositions of a Sigmoidal Function, *Mathematics of Control, Signals, and Systems*, vol. 2, no. 4, pp. 303–314, 1989, doi: 10.1007/BF02551274.

(2) B. M. Maschke and A. J. van der Schaft, Port-Controlled Hamiltonian Systems: Modelling Origins and System-Theoretic Properties, *IFAC Proceedings Volumes*, vol. 25, no. 13, pp. 282–288, 1992, doi: 10.1016/S1474-6670(17)52308-3.

(3) A. van der Schaft and D. Jeltsema, Port-Hamiltonian Systems Theory: An Introductory Overview, *Foundations and Trends in Systems and Control*, vol. 1, no. 2–3, pp. 173–378, 2014, doi: 10.1561/2600000002.

(4) S. Greydanus, M. Dzamba, and J. Yosinski, Hamiltonian Neural Networks, *Advances in Neural Information Processing Systems*, vol. 32, pp. 15353–15363, 2019, arXiv:1906.01563.

(5) M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations, *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019, doi: 10.1016/j.jcp.2018.10.045.

(6) M. Da Lio, D. Bortoluzzi, and G. P. Rosati Papini, Modelling Longitudinal Vehicle Dynamics with Neural Networks, *Vehicle System Dynamics*, vol. 58, no. 11, pp. 1675–1693, 2020, doi: 10.1080/00423114.2019.1638947.

(7) S. A. Desai, M. Mattheakis, D. Sondak, P. Protopapas, and S. J. Roberts, Port-Hamiltonian Neural Networks for Learning Explicit Time-Dependent Dynamical Systems, *Physical Review E*, vol. 104, no. 3, p. 034312, 2021, doi: 10.1103/PhysRevE.104.034312.

(8) G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, Physics-informed machine learning, *Nature Reviews Physics*, vol. 3, pp. 422–440, 2021, doi: 10.1038/s42254-021-00314-5.

(9) S. Eidnes, A. J. Stasik, C. Sterud, E. Bøhn, and S. Riemer-Sørensen, Pseudo-Hamiltonian Neural Networks with State-Dependent External Forces, *Physica D: Nonlinear Phenomena*, vol. 446, p. 133673, 2023, doi: 10.1016/j.physd.2023.133673.