Energy and AI 25 (2026) 100839

Contents lists available at ScienceDirect

Energy and AI

journal homepage: www.sciencedirect.com/journal/energy-and-ai

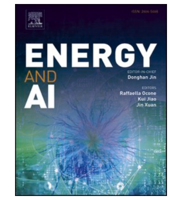

Learning-augmented model predictive control for energy management in hydrogen-battery hybrid trains: A comparative deep learning study

Yosef Berhan Jember [a,*], Massimo Santarelli [a], Giuseppe Santangelo [b], Clemente Mirko [b]

[a] Department of Energy, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129, Torino, Italy
[b] Alstom Ferroviaria S.p.A., Via O. Moreno, 23, 12038, Savigliano (CN), Italy

HIGHLIGHTS

- First four-model predictor comparison for hydrogen train energy management.
- The Echo State Network minimizes prediction error and battery charge drift.
- The Gated Recurrent Unit generalizes best, cutting hydrogen use by 13.4%.
- Long Short-Term Memory reduces fuel cell and battery stress the most.
- One-Dimensional Convolutional Network generalizes poorly on unseen routes.

GRAPHICAL ABSTRACT

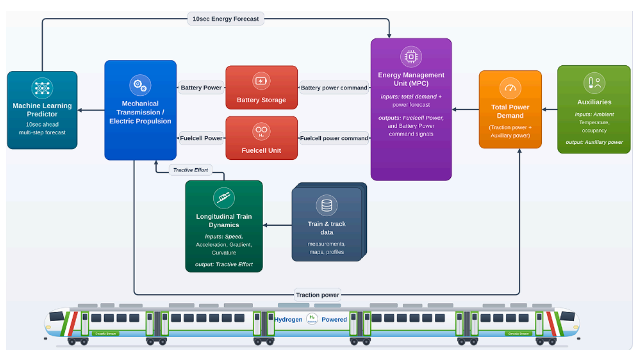

ARTICLE INFO

Keywords:
Hydrogen-battery hybrid trains
Energy management system
Model predictive control
Deep learning
Echo state network
Power demand forecasting

ABSTRACT

To decarbonize non-electrified rail networks, hydrogen-battery hybrid trains are emerging as promising alternatives to diesel traction. In such hybrid propulsion, the energy management system dynamically allocates power between the fuel cell and battery, making control design critical for hydrogen economy, charge sustainability, and component protection. Model predictive control combined with demand forecasting is a promising strategy, yet most prior studies adopt a single learning architecture and do not systematically assess how predictor choice affects closed-loop performance. This study presents a controlled comparison of four deep learning architectures: Long Short-Term Memory, Gated Recurrent Unit, One-Dimensional Convolutional Neural Network, and Echo State Network, as power-demand predictors within a unified model predictive control framework for hydrogen-train energy management. The control layer uses a zero-dimensional drivetrain model incorporating train dynamics, powertrain behavior, the fuel cell stack, and the battery system. The framework is developed using data from ALSTOM's Coradia Stream H train on the Brescia-Iseo-Edolo line and validated on two additional routes. Results reveal architecture-specific closed-loop trade-offs relative to standard model predictive control. The Echo State Network achieves the lowest prediction error and the smallest mean state-of-charge deviation, about 2.4%. Long Short-Term Memory yields the largest reductions in fuel-cell and battery loading stress, by 4.5 MW and 36.5 MW, respectively. Gated Recurrent Unit provides the highest hydrogen savings, up to 13.4%, with strong cross-route generalization, whereas the One-Dimensional Convolutional Neural Network overfits the training route

* Corresponding author.
E-mail address: yosef.jember@polito.it (Y.B. Jember).

https://doi.org/10.1016/j.egyai.2026.100839

Available online 3 July 2026

and degrades on unseen lines. The study provides deployment-oriented guidance for selecting learning-based predictors in hydrogen-train energy management.

1. Introduction

1.1. Background and motivation

The railway sector is facing increasing pressure to decarbonize while maintaining operational efficiency and reliability [1]. Hydrogen fuel cell technology has emerged as a viable solution for non-electric lines, providing zero-emission propulsion while offering a range comparable to diesel trains [2]. The Coradia Stream H is introduced as part of the H2IseO project on the Brescia-Iseo-Edolo line [3]. It becomes the first authorized commercial hydrogen train in the region [4], demonstrating the feasibility of this technology in a real-world operational environment.

However, the integration of hydrogen, fuel cells, and battery energy storage systems introduces complex energy management challenges in hybrid railway operation [5]. Unlike conventional diesel-electric trains, hydrogen-battery hybrid systems must optimize power distribution between fuel cells and batteries in real time to minimize hydrogen consumption, extend component lifetime, and ensure battery state-of-charge recovery at the end of the operational cycle. These objectives are often conflicting. For example, aggressive fuel cell load tracking strategies can minimize hydrogen consumption, but frequent output fluctuations accelerate fuel cell performance degradation [6,7]. Conversely, relaxed battery use can minimize fuel cell engagement, contributing to extending fuel cell lifetime, but it may compromise charge sustainability while increasing battery cycling and aging [8,9].

This challenge is further complicated in heavy-duty vehicles, including railway power demand, which varies with driving profile, route topography, operational schedules, and environmental conditions [5,10]. Existing rule-based energy management strategies are computationally efficient [5], but lack the adaptability required to optimize performance across various operating conditions. Model Predictive Control (MPC) provides a systematic multi-objective optimization framework capable of handling explicit constraints; however, its effectiveness strongly depends on accurate prediction of future power demand over the control horizon [11,12].

1.2. Literature review

Energy management strategies for hybrid electric vehicles have evolved from heuristic rule-based approaches to optimization-based methods. Early strategies relied on deterministic state machines and fuzzy logic controllers [13], which are robust but cannot guarantee optimality. Dynamic Programming (DP) can provide a globally optimal solution [14], but it requires complete prior knowledge of the driving cycle and is difficult to implement in real time due to its high computational complexity. The Pontryagin Minimum Principle (PMP) reduces computational burden but still depends on future driving information [15,16].

The Equivalent Consumption Minimization Strategy (ECMS) approximates the global optimal point through instantaneous optimization using an equivalent coefficient that relates electrical energy to fuel consumption [17]. Adaptive ECMS variants improve the robustness against cycle variability by adjusting this coefficient online based on driving conditions [18,19]. However, ECMS lacks explicit treatment of constraints related to component stress and charge sustenance, limiting its applicability to systems with strict performance degradation constraints [20].

MPC addresses these limitations by formulating energy management as a finite-horizon optimal control problem with explicit constraints on state and control variables [21]. Its receding horizon strategy enables real-time adaptation to changing conditions while maintaining computational flexibility [12,22]. Recent applications for fuel cell hybrid vehicles have demonstrated significant improvements in fuel economy and component protection compared to rule-based strategies. However, MPC performance remains fundamentally limited by the accuracy of predictions over the control horizon. The integration of machine learning with MPC has therefore emerged as a promising approach to improve predictive capability [23].

Deep learning architectures, especially Recurrent Neural Networks (RNNs), have shown superior performance in time-series prediction of energy systems. Recent works have applied learning-augmented MPC frameworks to fuel cell hybrid vehicles. In [24], a Bi-LSTM (Bidirectional Long Short-Term Memory) predictive-control Energy Management System (EMS) was proposed for a Fuel Cell Hybrid Electric Vehicle (FCHEV) to manage power distribution between fuel cells and energy storage systems. To reduce hydrogen consumption in Fuel Cell Hybrid Electric Buses (FCHEBs), an LMPC-LSTM (Learning MPC-LSTM) framework was developed [25] that integrates speed prediction. Xun et al. [26] introduced a reference-free Learning Model Predictive Controller (LMPC) for FCHEVs that optimizes power allocation through iterative learning. A cloud-assisted Economic Model Predictive Control (EMPC) strategy for fuel cell electric vehicles (FCEVs) that utilizes a learning-based terminal cost was also proposed [27] to enhance real-time efficiency. Jia et al. [28] established a learning-based model predictive EMS with health-aware control for FCHEBs using a double-layer Bi-LSTM speed predictor. An EMS integrating LSTM and DP within an MPC framework was presented to forecast power demand for FCHEVs [29]. Furthermore, a dual-model predictive control EMS based on a Wavelet Transform LSTM for online driving-condition recognition in FCEVs was proposed [30].

Beyond conventional predictive frameworks, a passenger-aware collaborative energy management strategy for fuel cell buses [31] integrates real-time occupant counts and gender distributions into a twin-delayed deep deterministic policy gradient algorithm, using a dual-source perception framework to coordinate power distribution and cabin thermal management. And a study in [32] proposed a decoupled safety supervision architecture for fuel cell vehicles that utilizes an independent safety-guided network to explicitly enforce lithium-ion battery thermal constraints, overcoming the mutual interference and reward-tuning difficulties inherent in traditional penalty-based methods.

1.3. Research gap and contributions

As described in the literature section, predictive and adaptive energy-management strategies for fuel-cell hybrid railway vehicles have been established through causal optimal-control, MPC-based, and forecast-informed formulations [5,11]. Related studies have also addressed predictive control and MPC-oriented hydrogen-train EMS [10,19]. In addition, it has been extended toward adaptive PMP-based MPC energy management [33]. Learning-based EMS in rail mobilities has also begun to emerge outside the strict MPC-predictor setting. Deep reinforcement learning has been used to develop railway EMS with degradation-aware objectives [34], while comparative AI-based EMS studies have also been reported [35]. The adoption of a learning-augmented MPC strategy addresses limitations found in conventional energy management frameworks. Rule-based or standard MPC strategies lack the real-time adaptability required to anticipate rapid, non-linear demand fluctuations across variable rolling stock operations, while global offline optimization methods are computationally prohibitive for online execution. Integrating machine-learning forecasting directly into the MPC optimization loop provides predictive demand awareness and improves power-splitting efficiency.

However, prior studies in learning-based EMS either adopt a single-predictor architecture, compare complete EMS strategies rather than alternative predictors within a common controller, or do not focus on hydrogen-battery railway applications within a unified MPC structure. Consequently, the literature still provides limited guidance on how predictor architecture affects closed-loop MPC energy-management behavior when the train model and control formulation are held fixed. This limitation is important because prediction quality directly shapes the controller's ability to anticipate future demand over the receding horizon, thereby affecting hydrogen consumption, battery charge-sustaining behavior, and fuel-cell loading stress. More importantly, prediction accuracy alone does not necessarily determine EMS quality. Different predictor architectures may yield distinct closed-loop control responses, even when embedded within the same MPC framework.

Moreover, the reservoir-computing family, represented by Echo State Networks, is especially underexplored in this context. Foundational work on reservoir computing established its potential for nonlinear sequence modeling and efficient recurrent computation [36, 37]. More recent evidence has shown strong Echo State Network (ESN) capability in time-series forecasting [38]. Nevertheless, the suitability of ESNs for learning-augmented MPC-based energy management in hydrogen-train applications has not yet been clearly established. To position the contributions more precisely, Fig. 1 summarizes representative prior work and highlights the narrower, unresolved question addressed here: how changes in predictor architecture affect closed-loop MPC energy management performance in hydrogen-battery hybrid trains within a common control framework.

As illustrated in Fig. 1, prior work has addressed predictive and learning-based railway and fuel-cell vehicle EMS domains, but a controlled comparison of multi-architecture predictors within a fixed hydrogen-train MPC loop, evaluated using closed-loop EMS outcomes, remains underexplored. The present work addresses this gap through three main contributions:

1. A systematic four-way comparison of Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), One-Dimensional Convolutional Neural Network (1D-CNN), and ESN predictors within a unified learning-augmented MPC framework for hydrogen-battery hybrid trains, enabling the effect of predictor architecture on closed-loop EMS behavior to be isolated under a fixed plant and control formulation.
2. A comprehensive closed-loop validation across three topographically diverse railway lines, establishing the relationship between predictor behavior and EMS performance metrics, including hydrogen consumption, component-degradation indicators, and charge-sustaining behavior.
3. The formulation of deployment-oriented design guidance linking operational priorities, such as fuel economy, component protection, State of Charge (SOC) regulation, and computational budget, to forecasting-architecture selection in learning-augmented MPC-based EMS.

1.4. Selection of deep learning architectures

Four architectures were considered for comparison on the basis of the following selection criteria: (i) it must represent a structurally distinct family of sequence modelling approaches, so that the comparison spans a meaningful range of inductive biases; (ii) the model should have documented precedent or a principled theoretical rationale for energy time-series prediction tasks; and (iii) compatibility with real-time MPC integration in line with railway traction system computational resources.

LSTM appears as an established recurrent baseline. Gated memory cells allow the network to selectively retain or discard information over arbitrarily long input sequences, addressing the vanishing-gradient issue that limits simpler recurrent architectures. Train driving cycles exhibit temporal structure at multiple scales, from fractional-second traction transients to multi-minute station-to-station segments, and LSTM is appropriate to capture this hierarchy. Its widespread adoption in hybrid vehicle EMS research makes it the default reference point for assessing other architectures.

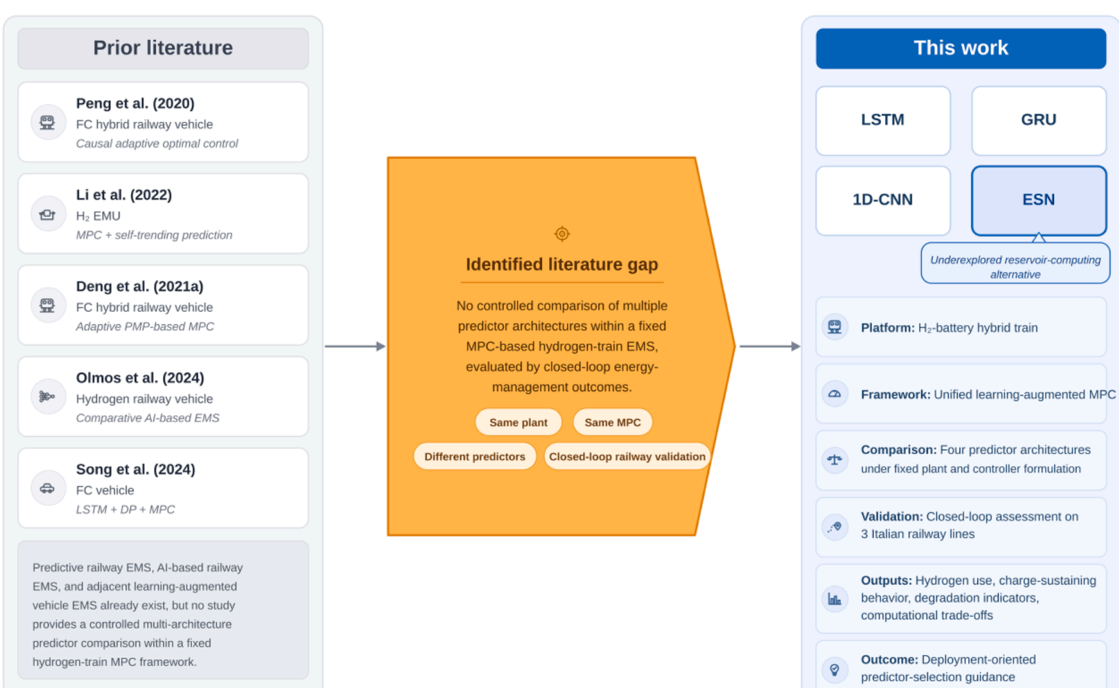

Fig. 1. Positioning of the proposed work relative to representative predictive and learning-based EMS studies for fuel-cell vehicles and hydrogen trains, and the gap it addresses.

GRU is a type of recurrent neural network and is a variant with fewer parameters. By consolidating the forgetting gate and the input gate into an integrated update mechanism, it operates with fewer parameters while maintaining time-series modeling capabilities equivalent to LSTM. This difference is important, especially in considering onboard deployment of the EMS. Given limited memory capacity and the need for regular retraining, an architecture that trains quickly and requires little storage without compromising predictive performance is advantageous.

1D CNN is selected to examine the hypothesis that local feature extraction, without any recurrent state, is sufficient for the 10-step prediction horizon demanded by the MPC. One-dimensional convolution identifies recurring local patterns and hierarchical features within a fixed receptive field and has been applied to short-term electrical load forecasting and predictive EMS in vehicles, yielding competitive results. Including CNNs provides a structurally non-recurrent contrast to LSTMs and GRUs, improving the interpretation of observed performance differences.

ESN represents the reservoir computing paradigm. A large, sparsely connected recurrent reservoir, initialized randomly and left untrained, projects input signals into a high-dimensional dynamical state space; only the linear readout weights are fitted by ridge regression. Eliminating backpropagation through time yields training times and parameter counts far below those of gradient-trained recurrent networks. ESN has shown particular strength in irregular, high-variance energy demand signals. Its inclusion here constitutes, to the authors' knowledge, the first evaluation of reservoir computing for hydrogen train power prediction.

The composition of this paper is as follows. Chapter 2 provides an overview of the hydrogen train system and operations. Chapter 3 describes the component modeling approach for train dynamics, fuel cells, and battery subsystems. Chapter 4 presents the deep learning architectures for power demand prediction. Chapter 5 details the MPC formulation and its integration with the trained predictors. Chapter 6 outlines the system settings, data sources, and validation methodology. Chapter 7 provides the results in terms of prediction accuracy and energy management performance. Chapter 8 discusses its interpretations, limitations, and future research directions, while Chapter 9 summarizes the key findings.

2. System description

Hydrogen technology is becoming increasingly important in the ecological transition, grounded in a model of sustainable development, decarbonization, renewable energy, and pollution reduction. In railway transportation, hydrogen trains represent a viable, environmentally friendly alternative to diesel trains operating on non-electrified lines.

A hydrogen train uses a hybrid architecture that combines fuel cells, traction batteries, traction inverters, hydrogen tanks, and electric motors. Fig. 2 shows the schematic configuration of this train with indication of its power and energy conversion units.

In this hybrid architecture, energy and power management play a critical role in train performance, as demonstrated by the diagram in Fig. 3. The train's energy and power performance can be assessed when key operational constraints are defined:

- Mission profile (track gradient, curves, speed limits, climate, timetable, passenger load, component health).
- Battery state of charge and power flow, which deeply affect battery aging and lifetime.

Available onboard hydrogen, which determines mission autonomy.

The main objective of the EMS is to minimize total energy consumption over the mission, reduce the total cost of ownership, and maximize performance and component lifetime. This is achieved through optimization processes that integrate mission data, component characteristics, and real-time measurements.

The case study hydrogen train performance characteristics are shown in Table 1 from a prior study that aimed to develop a modeling and sizing optimization approach to evaluate the operational efficiency and economic viability of these hybrid fuel cell-battery trains.

The motivation behind the present research is to improve energy optimization, enhance operational efficiency, and accelerate innovation in green mobility, contributing to a more sustainable future.

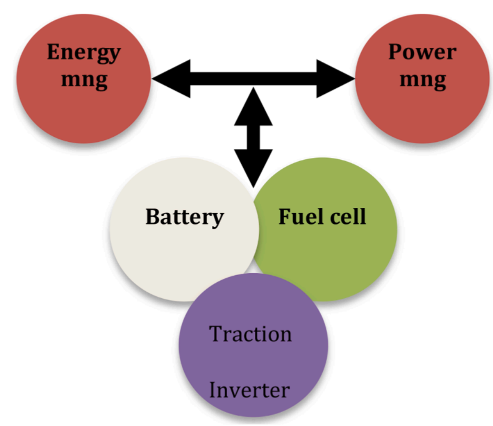

Fig. 3. Conceptual representation of power flow and energy management (mng) coordination in the hybrid powertrain.

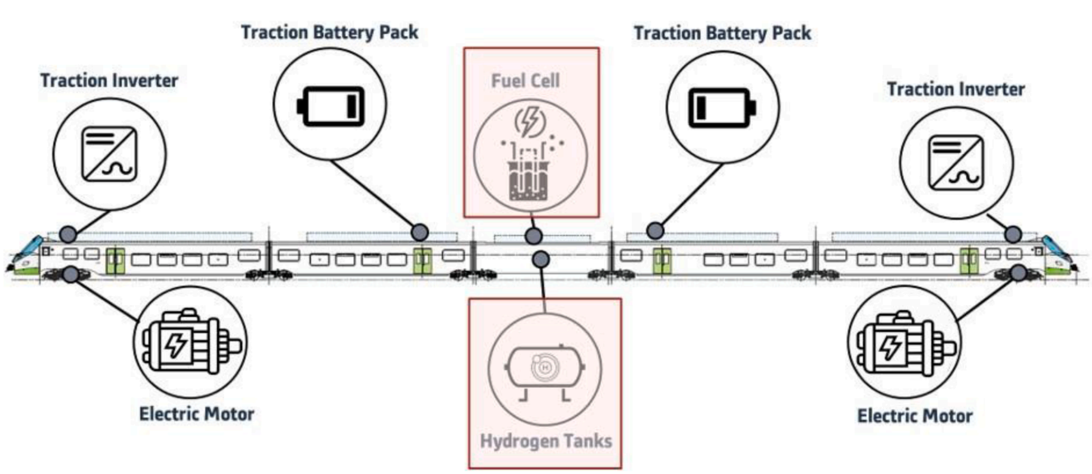

Fig. 2. Powertrain architecture of the studied hydrogen-battery hybrid train, with its main components indicated.

Table 1
Main specifications of the case-study hydrogen train [39].

Characteristics	Value
Maximum weight [t]	216
Maximum acceleration-deceleration [m/s²]	± 0.8
Maximum traction power [kW]	1170
Train autonomy [km]	600
Total seats [-]	260
Rotary allowance λ [%]	7.0

3. Component modeling

The system modeling approach adopted in this work is based on a Zero-Dimensional (0D) simulation framework that represents longitudinal train dynamics, powertrain behavior, auxiliary unit, fuel cell characteristics, and battery system [39]. This level of abstraction strikes a balance between computational efficiency and the accuracy required to develop energy management strategies.

The overall architecture of the proposed framework, depicted in Fig. 4, integrates train dynamics, a machine-learning model for power prediction, and an MPC-based energy-management strategy. MATLAB/Simulink and Python are used as the main simulation development environments.

3.1. Train longitudinal dynamics

The train's motion is governed by Newton's second law, balancing tractive effort with resistive forces [40]. The equation is expressed as:

$$M_{train}\frac{dv}{dt} = F_{traction} - F_{resistance} - F_{grade} - F_{curve} \tag{1}$$

Where M_{train} is the total train mass, v is velocity, $F_{traction}$ is the tractive force from the propulsion system, and the resistance force terms represent the total tractive resistance $F_{resistance}$, grade resistance F_{grade}, and curve resistance F_{curve}.

The total tractive resistive force is typically modeled using the Davis equation [41]:

$$F_{resistance} = A + Bv + Cv^2 \tag{2}$$

Where the coefficients A, B, and C represent rolling resistance, mechanical friction, and aerodynamic drag, respectively. For the Coradia Stream H train, these coefficients were obtained from the manufacturer's specifications and verified through simulations.

The grade resistance depends on the track alignment angle θ, while the radius along the track r determines the curve resistance, and it is computed based on railway design practice:

$$F_{grade} = M_{train} g \sin(\theta) \tag{3}$$

$$F_{cur} = \frac{800 M_{train}\, g}{r} \tag{4}$$

Where g is the gravitational acceleration and r the radius of curvature.

The power demand at the wheel is given by:

$$P_{demand}(t) = F_{traction}(t)v(t)/\eta_{transmission} + P_{aux}(t) \tag{5}$$

This power demand must be supplied by the fuel cell and battery system and constitutes the input to the energy management controller.

$P_{aux}(t)$ represents auxiliary loads, such as Heating, Ventilation, and Air-Conditioning (HVAC), lighting, and control systems, which are modeled based on their power requirements with temperature variations.

Where $\eta_{transmission}$ dynamically models the combined efficiency of the propulsion powertrain. While a constant efficiency accounts for steady mechanical losses in the transmission gear.

The electric drive components are modeled dynamically to capture operational non-linearities. The speed-dependent efficiency profile of the traction motors was characterized via the JMAG simulation environment, leveraging manufacturer nameplate data, available technical specifications, and traction literature to approximate the non-linear curve shown in Fig. 5.

In power electronics, while traction converters are generally highly efficient, their performance varies nonlinearly with operating load and power. To capture these dynamics, a modified traction converter model

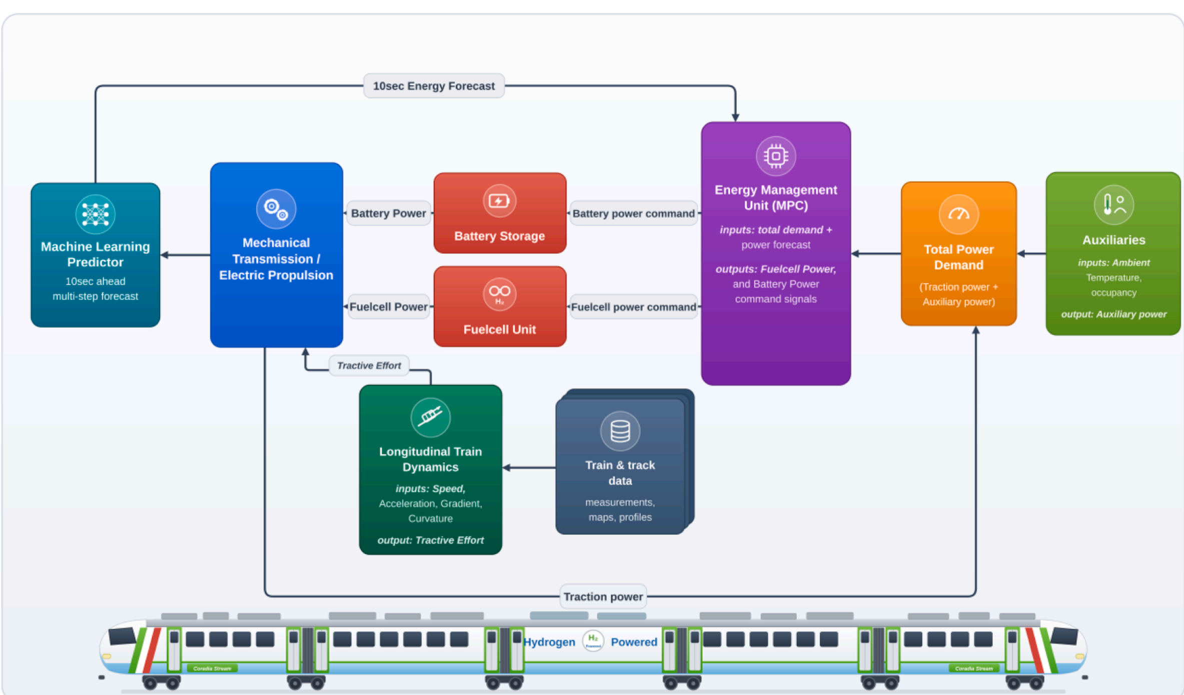

Fig. 4. Block diagram of the proposed learning-augmented MPC framework and its coupling to the train physical system, with the main signal exchanges between subsystems indicated.

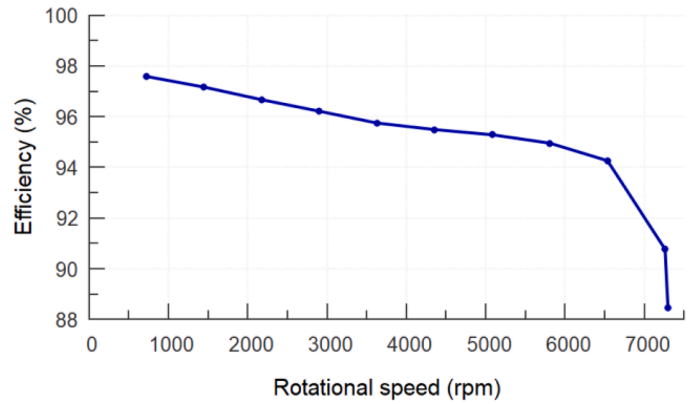

Fig. 5. Speed-efficiency curve of the traction motor, approximated from JMAG simulation, showing the decline in efficiency with increasing rotational speed.

was established by cross-referencing validated electrical drive literature, including performance evaluations of Silicon Insulated-Gate Bipolar Transistors (Si IGBTs), silicon carbide Metal-Oxide-Semiconductor Field-Effect Transistors (MOSFETs), and hybrid switch topologies. These benchmarked literature curves are compiled in Fig. 6, where the Si IGBT-based traction drive profile represents the specific non-linear behavior adopted for the inverter model in this study.

3.2. Fuel cell system

Fuel cells convert hydrogen gas into electrical energy through electrochemical reactions [44]. In this modeling approach, a zero-dimensional quasi-static model based on polarization curves was adopted, which provides sufficient accuracy for energy management studies while maintaining computational efficiency [45,46].

The fuel cell voltage is modeled as:

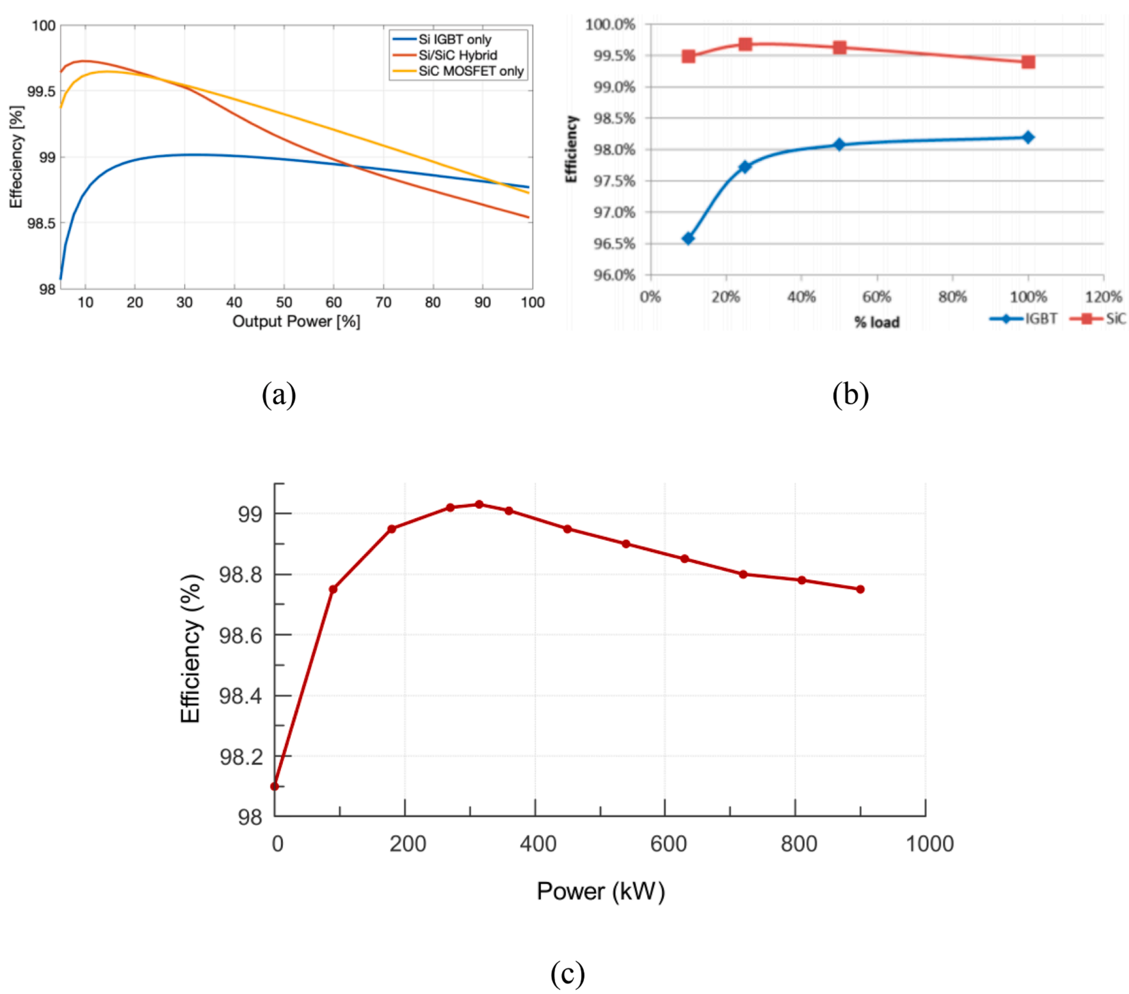

Fig. 6. Comparative traction-inverter efficiency profiles across switch topologies from the literature: (a) Si IGBT, Si/SiC hybrid, and SiC MOSFET traction drives [42]; (b) IGBT vs. SiC traction inverter [43], and (c) the Si IGBT-based profile adapted for the inverter model in this study. All panels plot efficiency against output power or load.

$$V_{FC} = E_{Nernst} - V_{act} - V_{ohm} - V_{conc} \tag{6}$$

Where E_{Nernst} is the thermodynamic open-circuit voltage, and V_{act}, V_{ohm}, and V_{conc} represent activation, ohmic, and concentration overpotentials, respectively [44,47].

For the fuel cell stack representative of ALSTOM hydrogen train designs, polarization curves were derived from primary data and adjusted using literature sources [3]. The electrical power output is given by:

$$P_{FC} = V_{FC} I_{FC} \tag{7}$$

Where I_{FC} is the stack current. The hydrogen consumption rate is calculated using a corrected form of Faraday's law:

$$\dot{m}_{H_2} = \frac{M_{H_2} I_{FC}}{2F\eta_{FC}} \tag{8}$$

Where M_{H_2} is the molar mass of hydrogen, F is Faraday's constant, and η_{FC} is the Faradaic efficiency accounting for parasitic losses [44,45].

The performance degradation of fuel cells is mainly caused by load cycling and operation at extreme power levels [7,48]. In this study, Fuel Cell (FC) stress is quantified as the cumulative absolute variation in output power:

$$\text{Stress}_{FC} = \sum_{k=1}^{N_{trip}} |P_{FC}(k) - P_{FC}(k-1)| \tag{9}$$

k is the timestep index, and N_{trip} stands for the total number of simulation time steps for the trip. This indicator correlates with membrane degradation, catalyst dissolution, and carbon corrosion observed in accelerated aging tests [48,49]. Lower stress values imply smoother power profiles, which lead to longer fuel cell lifetime [7,50].

3.3. Battery energy storage system

The battery is modeled using the SOC dynamics:

$$SOC(t+1) = SOC(t) - \frac{\eta_{bat}^{\pm 1} P_{bat}(t)\Delta t}{Q_{bat} V_{nom}} \tag{10}$$

Where Q_{bat} is the battery capacity, V_{nom} is the nominal voltage, and η_{bat} is charge/discharge efficiency. Battery capacity and power ratings specific to the case study hydrogen train are considered in this work.

Performance degradation of the battery is affected by various factors such as cycle depth, C-rate, temperature, and SOC operating range [51,52]. Battery (bat) stress is calculated in a similar manner to the fuel cell:

$$\text{Stress}_{bat} = \sum_{k=1}^{N_{trip}} |P_{bat}(k) - P_{bat}(k-1)| \tag{11}$$

This metric captures cumulative power cycling that causes capacity reduction and impedance increase [9].

In addition, the charge sustaining is tracked through the absolute SOC deviation:

$$\Delta SOC = |SOC_{end} - SOC_{start}| \tag{12}$$

Charge-sustaining operation (ΔSOC ≈ 0) ensures that the battery returns to its initial state after a round trip by utilizing regenerative braking power, preventing gradual discharge or overcharging.

4. Deep learning architectures for power demand prediction

4.1. Problem formulation

The power demand forecasting task is formulated as a univariate time-series prediction problem. Given the past power demand $\{P_{demand}(t-L), \ldots, P_{demand}(t-1)\}$ observed over a lookback window of length L, the objective is to predict the future power demand $\{\widehat{P}_{demand}(t), \ldots, \widehat{P}_{demand}(t+H-1)\}$ over a prediction horizon of length H [23]. In this work, the lookback window L is set to 60 seconds, and the prediction horizon H to a 10second timestep, providing sufficient historical context and prediction horizon for effective MPC implementation. The model adopts a direct multi-horizon prediction strategy, i.e., a multi-input multi-output formulation, which directly outputs the entire prediction horizon as a vector [53].

The univariate configuration was selected to optimize the balance between computational speed and model generalizability. Preliminary development evaluations demonstrated that incorporating multi-feature operational or topographic inputs did not yield a statistically significant reduction in prediction errors. Consequently, the streamlined univariate approach was maintained to minimize model complexity and eliminate dependency on external sensor telemetry arrays.

4.2. Long short-term memory

The LSTM network addresses the vanishing gradient problem of conventional RNNs through gated memory cells that selectively retain or discard information over long sequences [54]. The update equations of the LSTM cell are given by:

$$\begin{aligned}
f_t &= \sigma(W_f[h_{t-1}, x_t] + b_f) \\
i_t &= \sigma(W_i[h_{t-1}, x_t] + b_i) \\
\widetilde{C}_t &= \tanh(W_C[h_{t-1}, x_t] + b_C) \\
C_t &= f_t \odot C_{t-1} + i_t \odot \widetilde{C}_t \\
o_t &= \sigma(W_o[h_{t-1}, x_t] + b_o) \\
h_t &= o_t \odot \tanh(C_t)
\end{aligned} \tag{13}$$

where f_t, i_t, and o_t denote the forget, input, and output gates, respectively; x_t input vector at timestep t, C_t cell state, $\widetilde{C}_t$ candidate cell state, C_{t-1} previous cell state, h_t the hidden state, h_{t-1} previous hidden state, σ sigmoid function, (W_f, W_i, W_C, W_o) weight matrices, (b_f, b_i, b_C, b_o) bias vectors, tanh (hyperbolic tangent) activation function, and $\odot$ indicate element-wise multiplication.

LSTM networks can learn long-term dependencies [55] that are important for capturing periodic patterns in rail power demand. The implemented model consists of two stacked LSTM layers with dropout regularization, followed by a dense layer outputting a 10-step prediction.

4.3. Gated recurrent unit

GRU simplifies the LSTM architecture by integrating the forget and input gates into a single update gate and merging the cell state with the hidden state [56,57]. The GRU update equations are:

$$\begin{aligned}
z_t &= \sigma(W_z[h_{t-1}, x_t] + b_z) \\
r_t &= \sigma(W_r[h_{t-1}, x_t] + b_r) \\
\widetilde{h}_t &= tanh(W_h[r_t \odot h_{t-1}, x_t] + b_h) \\
h_t &= (1 - z_t) \odot h_{t-1} + z_t \odot \widetilde{h}_t
\end{aligned} \tag{14}$$

where z_t represents the update gate and r_t the reset gate. Compared to LSTM, GRU has fewer parameters, which often leads to faster training and lower risk of overfitting, while maintaining comparable performance in many sequence modeling tasks.

The GRU model consists of two GRU layers with dropout and batch normalization, followed by a dense layer that outputs the prediction steps.

4.4. Convolutional neural network

Although CNNs are traditionally associated with spatial data, one-dimensional convolution has proven effective for time series modeling by learning local patterns and hierarchical features [58,59]. The 1D convolution operation is defined as:

$$y_t = \sum_{k=0}^{K-1} w_k x_{t-k} + b \tag{15}$$

Where w_k are learnable filter weights, K denotes the kernel size, and b is the bias term. Stacking multiple convolutional layers with increasing dilation rates allows the network to capture patterns at different time scales [58].

The adopted architecture consists of two 1D convolutional layers with Rectified Linear Unit (ReLU) activation functions, followed by dropout regularization, flattening, and a final dense layer that produces a 10-step prediction. This hierarchical feature extraction method enables CNN to identify temporal patterns [59] in power demand profiles, such as acceleration, constant-speed, and braking phases.

4.5. Echo state network

ESN belongs to the reservoir computing paradigm, in which a large, fixed, random recurrent network (the reservoir) projects input signals into a high-dimensional space, and only the output weights are trained [36,37]. The reservoir dynamics are described by:

$$x(t) = (1-\alpha)x(t-1) + \alpha tanh(W_{in}u(t) + W_{res}x(t-1)) \tag{16}$$

where $x(t)$ is the reservoir state, $u(t)$ is the input, W_{in} and W_{res} are the fixed random input and reservoir weight matrices, and α is the leaking rate. The output is computed as:

$$y(t) = W_{out}x(t) \tag{17}$$

where W_{out} is trained by linear regression, typically ridge regression [37,60]. This training approach is computationally efficient, requiring only solving a linear system rather than iterative gradient descent [36].

The ESN is implemented using standard hyperparameter settings (reservoir size, spectral radius, input scaling, leaking rate, and sparsity) to balance dynamics and stability [37,60]. To guarantee network stability and satisfy the Echo State Property, the stochastic reservoir weight matrix W_{res} is normalized by its maximum absolute eigenvalue and scaled to a fixed spectral radius $\alpha < 1$.

5. Model predictive control framework

5.1. MPC formulation

The energy management problem is formulated as a finite-horizon optimal control problem solved at each timestep t [21]. The objective function is defined as:

$$J = \sum_{i=0}^{p-1} \Big[W_{Soc}\left(SoC(i+1) - SoC^{ref}(i+1)\right)^2 + W_{P_{bat}}P_{bat}(i)^2 + W_{P_{fc}}P_{fc}(i)^2 + W_{\Delta P_{bat}}\Delta P_{bat}(i)^2 + W_{\Delta P_{fc}}\Delta P_{fc}(i)^2 \Big] + pen \tag{18}$$

Subject to constraints:

$$P_{dem}(i) = P_{fc}(i) + P_{bat}(i),\ i = 0,\ \dots,\ p-1 \tag{19}$$

$$SoC(i+1) = f(SoC(i),\ P_{bat}(i)) \tag{20}$$

$$P_{bat,min} \le P_{bat}(i) \le P_{bat,max} \tag{21}$$

$$P_{fc,min} \le P_{fc}(i) \le P_{fc,max} \tag{22}$$

$$SoC_{min} \le SoC(i) \le SoC_{max} \tag{23}$$

$$\Delta P_{bat,min} \le \Delta P_{bat}(i) \le \Delta P_{bat,max} \tag{24}$$

$$\Delta P_{fc,min} \le \Delta P_{fc}(i) \le \Delta P_{fc,max} \tag{25}$$

where p is the prediction horizon, SoC is the state of charge, SoC^{ref} is the reference state of charge, P represents power, and ΔP denotes the change in power. The subscripts *bat*, *fc*, *dem*, identify the battery, fuel cell, and demand, respectively, while *min* and *max* indicate the operational limits. The index (i) refers to the current interval and $(i+1)$ to the next predicted step.

The objective function balances several conflicting goals through adjustable weighting factors (W_{Soc}, $W_{P_{bat}}$, $W_{P_{fc}}$, $W_{\Delta P_{bat}}$, $W_{\Delta P_{fc}}$). These include tracking the reference SoC trajectory to maintain battery charge, limiting component degradation by penalizing excessive battery and fuel cell power usage, reducing hydrogen consumption, and mitigating rapid changes in power output to limit component stress [11,12]. A penalty term (*pen*) is also included to relax constraints and keep feasibility under predictive uncertainty.

The power-balance constraint ensures that the required traction energy is met, while the discrete-time SoC dynamics govern the battery's energy evolution. Constraints on component power limits, SoC bounds, and power rate variations define the optimal operating range and protect against excessive performance degradation [7,8].

5.2. Quadratic programming reformulation

The fuel cell power trajectory is treated as a decision variable, while the battery power is determined from the power balance constraint. Quadratic penalties on fuel cell and battery power encourage operation in high-efficiency regions, thereby indirectly reducing hydrogen use and energy losses. Similarly, quadratic terms on power changes penalize rapid transients, effectively estimating component stress and supporting longer lifespans [11,29].

The resulting optimization problem can be expressed in the standard Quadratic Programming (QP) form:

$$\begin{aligned} &\min_x \frac{1}{2}x^T Q x + q^T x \\ &\text{subject to } Gx \le h \\ &\qquad Ax = b \end{aligned} \tag{26}$$

The decision vector is $x = \left[P_{fc}(t),\ \cdots,\ P_{fc}(t+p-1),\ SoC(t+1),\ \cdots,\ SoC(t+p)\ \right]^T$, Q is the Hessian matrix of the quadratic cost function, and q, G, h, A, and b represent the linear cost components and the boundaries for inequality and equality constraints, respectively.

This QP is solved in real time using the active-set method. Within the receding-horizon framework, solutions from the previous horizon are used as warm starts, enabling fast convergence and low computational latency.

5.3. Integration with deep learning predictors

The MPC controller requires future power demands $\{P_{demand}(t), \dots, P_{demand}(t+H-1)\}$ across the prediction horizon [23]. In the learning-augmented framework, these values are provided by trained deep learning models. Although predicted values are used for power splitting, the train's energy demand is satisfied in all learning-augmented cases.

$$[\widehat{P}_{demand}(t), \dots, \widehat{P}_{demand}(t+H-1)] = f_{DL}([P_{demand}(t-L), \dots, P_{demand}(t-1)]) \tag{27}$$

Where f_{DL} represents the LSTM, GRU, CNN, or ESN predictor.

At each timestep, the workflow proceeds as follows:

1. Collect past power demand over the lookup window.
2. Normalize the inputs using training statistics,
3. Generate predictions with deep learning models, and inversely normalize the predictions to their original scales.
4. Formulate and solve the MPC optimization problem,

5. Apply the first control action, corresponding to the fuel cell output setting point
6. Repeat the process by moving to the next timestep.

This receding-horizon strategy enables continuous adaptation to system behavior and corrects prediction errors at each step via feedback. The integration of the predictor and MPC combines the pattern-recognition capabilities of deep learning with the constraint-handling and effective optimization structure of model-based control [28,29].

5.4. *Baseline strategy*

For comparison purposes, a standard MPC controller without learning-augmented prediction is implemented. The baseline assumes constant power demand over the prediction horizon:

$$\widehat{P}_{\text{demand}}\ (k) = P_{\text{demand}}\ (t-1), \quad k = t, \ldots, t+H-1 \tag{28}$$

This persistence forecast represents a simple baseline prediction strategy in predictive energy-management frameworks that rely on future-demand forecasts [23]. Comparing the learning-augmented MPC against this baseline isolates the value added by deep learning prediction.

6. System setup and validation methodology

6.1. *Data sources and route characteristics*

The map in Fig. 7 illustrates the geographical locations, distances, and line geometries of the railway routes used in the investigation of the proposed EMS.

The study adopts typical hydrogen train parameters [39], and combine existing line data with simulated mission profiles from three distinct railway routes. These routes provide diverse topological and operational characteristics for model validation.

6.1.1. *Training route: Brescia-Iseo-Edolo*

The Brescia-Iseo-Edolo line serves as the primary training route. The Coradia Stream H hydrogen train is currently undergoing a test run on this line and is expected to enter official operation in 2026. This 103 km non-electrified rail line spans varied terrain with an elevation change from 130 m to 670 m above sea level, 34 intermediate stops, mixed urban and rural operating environments, and a 2.5-hour driving time.

6.1.2. *Validation route 1: Limone-Ventimiglia*

This 70 km mountainous route includes diverse topography: a sea-level to 1100 m elevation change, 9 intermediate stops, predominantly tunneled coastal terrain, and a driving time of 3.5 hours.

6.1.3. *Validation route 2: Firenze - Faenza*

This 100 km route also features distinct operating conditions: an elevation change from 50 m to 615 m above sea level, 17 intermediate stops, flat agricultural terrain, and a driving time of 3.7 hours.

By considering these three different railway lines, the generalization ability and route-dependent performance characteristics of the proposed models can be effectively evaluated.

6.2. *Model training and implementation*

The models are trained on data from the company's high-fidelity commercial simulation tool for the Brescia-Iseo-Edolo route. The development process consists of the following steps:

1. Data preprocessing and normalization are performed. The power demand profile is min-max scaled to the range [0,1] using statistics derived from the training dataset.
2. An input sequence is generated using a 60-second sliding window, and the corresponding 10-step future power demand is set as a target value.
3. Dataset segmentation is applied in chronological order, with the first 80% of sequences being used for learning and the remaining 20% for validation and testing.

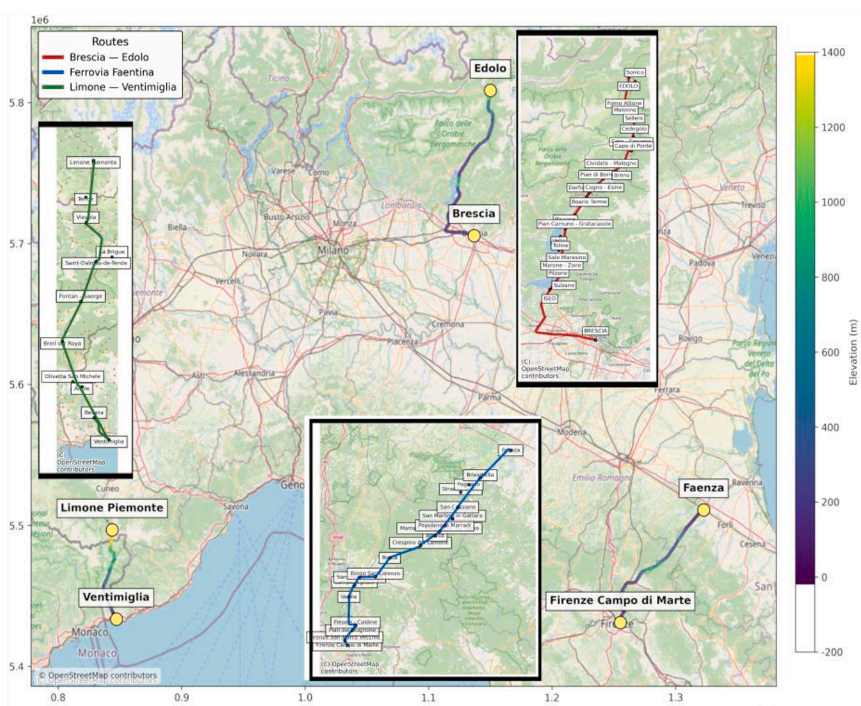

Fig. 7. The three railway lines used as case studies: Brescia-Iseo-Edolo, Firenze-Faenza, and Limone-Ventimiglia, showing route geometry, station locations, and track elevation.

4. Deep learning models (LSTM, GRU, CNN) are trained end-to-end using Adam optimizer and the mean square error loss function. In contrast, ESN uses a fixed reservoir configuration, while the linear readout layer is trained by ridge regression with L2 regularization [37,60].
5. Finally, the trained model is evaluated on the Brescia-Iseo-Edolo test set data, and further analyzed the generalization performance through MPC-EMS on the other two routes.

All deep learning models are implemented using TensorFlow/Keras, while ESN uses custom NumPy/scikit-learn implementations to ensure consistency and reproducibility with established reservoir computing practices.

6.3. Performance evaluation and validation

The validation process consists of three steps:

Step 1: Assessing prediction accuracy

Each trained model performs a 10-step-ahead prediction on a test dataset from the Brescia-Iseo-Edolo line. To analyze performance characteristics across the prediction horizon, the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and coefficient of determination (R^2) are computed for each prediction timestep.

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}|\hat{y}_i - y_i| \tag{29}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i - y_i)^2} \tag{30}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2} \tag{31}$$

Step 2: Training route energy management

The learning-augmented MPC strategies (LSTM-MPC, GRU-MPC, CNN-MPC, ESN-MPC) and the baseline MPC are simulated for the Brescia-Iseo-Edolo round-trip route. Energy management performance metrics are then calculated and compared, considering hydrogen consumption, component stress, and SOC difference, using Eqs. (8-12).

Step 3: Validation routes generalization

The trained models with MPC controllers are simulated on the Limone-Ventimiglia and Firenze-Faenza lines to evaluate their performance. Energy management metrics are computed to analyze generalization capabilities and route-dependent performance characteristics.

All simulations use the same MPC parameters to isolate the effect of prediction quality on energy management performance. The weights ($W_{Soc}, W_{P_{bat}}, W_{P_{fc}}, W_{\Delta P_{bat}}, W_{\Delta P_{fc}}$) of the MPC objective function determines the trade-off between hydrogen consumption, component stress, and charge sustaining [21]. In this study, balanced weights were used to ensure an unbiased comparison. The weights were calibrated through an iterative heuristic tuning process to ensure sustained full-battery charge after a complete round trip, while balancing hydrogen consumption and component stress limits. The parameter-tweaking process began with uniform scaling of all cost terms, followed by iterative adjustments to identify the weight combination that maximizes state-of-charge recovery while maintaining hydrogen (H_2) fuel economy. However, in practical deployments, these weights can be adjusted based on operational priorities and economic considerations [11,12]. For instance, increasing the weight on hydrogen consumption prioritizes fuel economy, whereas increasing the weight on stress-related factors prioritizes component longevity. Systematic weight tuning using multi-objective optimization techniques can help identify optimal settings for specific operational contexts.

7. Result

7.1. Prediction performance analysis

Fig. 8 presents the prediction accuracy metrics across the 10-step forecast horizon for all four deep learning architectures evaluated on the Brescia-Iseo-Edolo test set. MAE and RMSE are expressed in kW, whereas R^2 is reported in its standard unit scale.

ESN showed good short-term prediction accuracy, achieving the lowest MAE values (approximately 18-49 kW) and the least RMSE values over the first three timesteps. This strength stems from the fact that high-dimensional random projections of the reservoir effectively capture the instantaneous temporal dynamics. Across the initial timesteps, the 1D CNN exhibits moderate performance (MAE ≈ 43-91 kW) and exploits hierarchical feature extraction for local patterns. LSTM and GRU yield relatively higher errors (MAE ≈ 56-104 kW), indicating lower sensitivity to very short-term fluctuations compared to ESN and CNN-based models. All models score excellent goodness of fit, with R^2 values exceeding 0.89 at timestep 1, gradually decreasing to about 0.73-0.75 by timestep 3.

As the prediction horizon increases, all models showed the expected degradation in performance, with MAE and RMSE increasing and R^2 decreasing. ESN maintained the lowest overall error (MAE ≈ 87-103 kW, RMSE ≈ 249-276 kW), but the relative performance decline may reflect the absence of trainable recurrent gates for preserving long-range temporal information [37,60]. LSTM and GRU followed (MAE ≈ 136-170 kW, RMSE ≈ 324-360 kW), benefiting from gate architectures providing moderate long-term robustness [54,56]. CNN experienced the largest performance degradation (MAE ≈ 158-203 kW, RMSE ≈ 323-363 kW), suggesting weaker handling of the longer temporal structure of the forecasting task [58,59]. The R^2 values for all models fell from >0.89 in timestep 1 to about 0.34-0.36 in timestep 10, implying increased uncertainty in capturing magnitude and variability across long-term horizons.

Fig. 9(a) illustrates the model performance during a high-power traction interval in the Brescia-Iseo-Edolo test cycle, showing a sample of the multi-stage power-demand forecast over a 50-second window. GRU and ESN tracks actual profiles most closely, while LSTM struggles to preserve peak magnitude and shapes. CNN shows larger deviations over the continuous high-power interval. Fig. 9(b) presents a representative prediction interval during a regenerative braking phase. ESN again achieves the closest match to the actual power in the short term, while LSTM and GRU show distinct deviations during the transition period, though they capture the overall trend. CNN and ESN show good short-term accuracy but display increasing divergence during the recovery phase.

These predictive properties provide implications for the MPC performance calculated for the overall driving simulation. Good accuracy across the ESN horizon, especially in the short-term phase, could enable precise control decisions and potentially support robust multi-step optimization.

However, as detailed in the following sections, the LSTM and GRU architectures exhibit strong closed-loop MPC performance despite higher prediction errors, which may be related to their ability to learn persistent patterns of power demand [28,29]. While CNN's intermediate prediction performance suggests moderate potential, the actual MPC effects appear limited and inconsistent across routes.

7.2. Training route performance (Brescia-Iseo-Edolo)

Table 2 presents energy management performance on the Brescia-

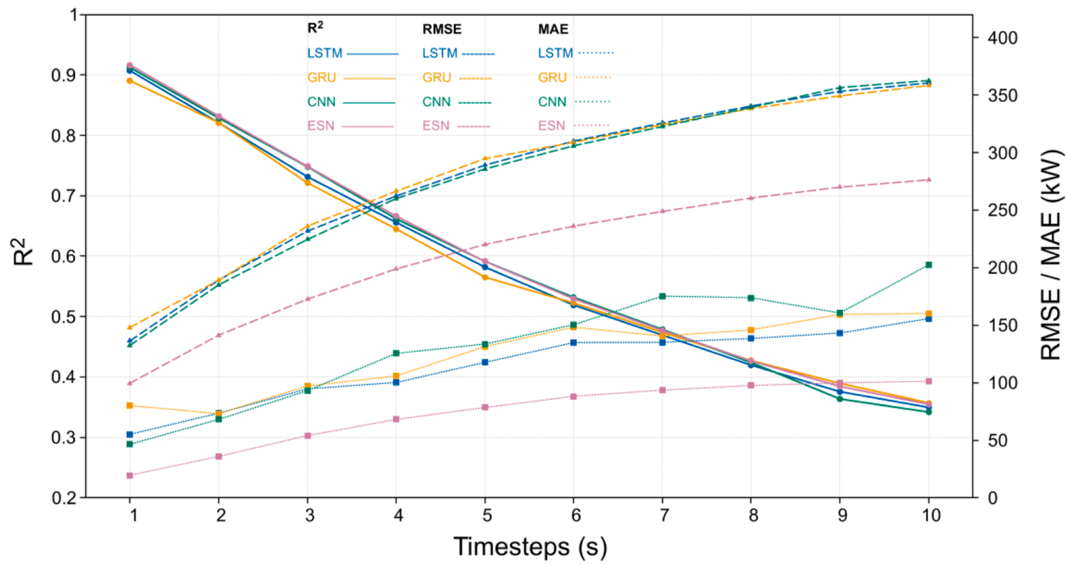

Fig. 8. Prediction performance of the four models across the 10-step horizon: R^2 (left axis) and RMSE/MAE (right axis, kW).

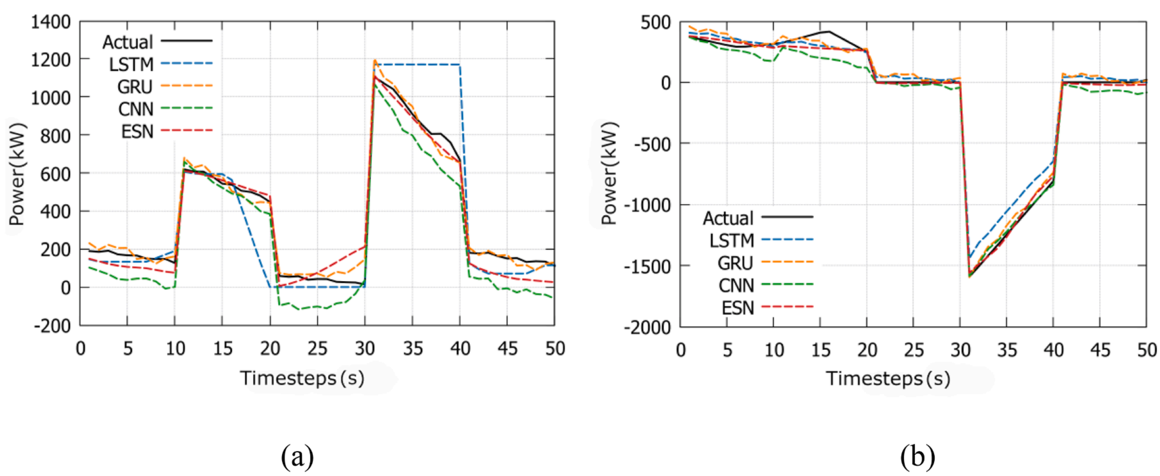

Fig. 9. Sample actual vs. predicted power for the four models: (a) a traction interval and (b) a regenerative-braking interval.

Table 2
Energy management performance on the Brescia-Iseo-Edolo route, comparing the standard MPC with the four learning-augmented strategies.

Strategy	H_2 Reduction (%)	FC Stress (MW)	Battery Stress (MW)	Δ SOC (%)
Standard MPC	Baseline (0)	60.0	238.5	0.12
LSTM-MPC	8.4	55.5	202.0	0.00
GRU-MPC	2.3	57.3	208.0	0.00
CNN-MPC	20.3	67.5	214.0	0.00
ESN-MPC	6.0	56.0	208.0	0.00

Iseo-Edolo training route, comparing the baseline MPC with four learning-augmented strategies.

The reductions in hydrogen consumption (kg) and battery energy usage (kWh) for the learning-augmented MPC cases relative to the reference MPC are also shown in Fig. 10.

7.2.1. Hydrogen consumption

CNN-MPC achieved the highest reduction in hydrogen consumption, at 20.3%, on the training line. This suggests the predictor was effective in supporting fuel-economy-oriented control actions under the training-route conditions. LSTM-MPC recorded a decrease of 8.4%, and ESN-MPC and GRU-MPC a decrease of 6.0% and 2.3%, respectively. The relatively conservative performance of GRU-MPC is consistent with high prediction error and indicates that optimizing hydrogen consumption on the training line can be sensitive to overall prediction accuracy or to differences in learned patterns [56,57]. Nonetheless, all learning-augmented strategies outperform the baseline MPC, confirming that even moderate predictive improvements lead to measurable fuel savings.

7.2.2. Fuel cell stress

LSTM-MPC achieved the lowest fuel cell stress at 55.5 MW, a 7.5% reduction from the 60.0 MW baseline MPC. This result is consistent with the LSTM-based predictor promoting a smoother fuel-cell power profile with fewer abrupt load changes. Reduced cumulative power fluctuations support longer membrane life and reduced catalyst dissolution, with possible cost benefits [48,50]. ESN-MPC and GRU-MPC achieved similar stress levels of 56.0 MW and 57.3 MW, respectively, showing a moderate decrease in stress. It should be noted that CNN-MPC increased fuel cell stress to 67.5 MW, representing a 12.5% increase from the baseline. This deterioration is due to CNN's more frequent adjustments to fuel cell output to take advantage of predicted demand fluctuations, while

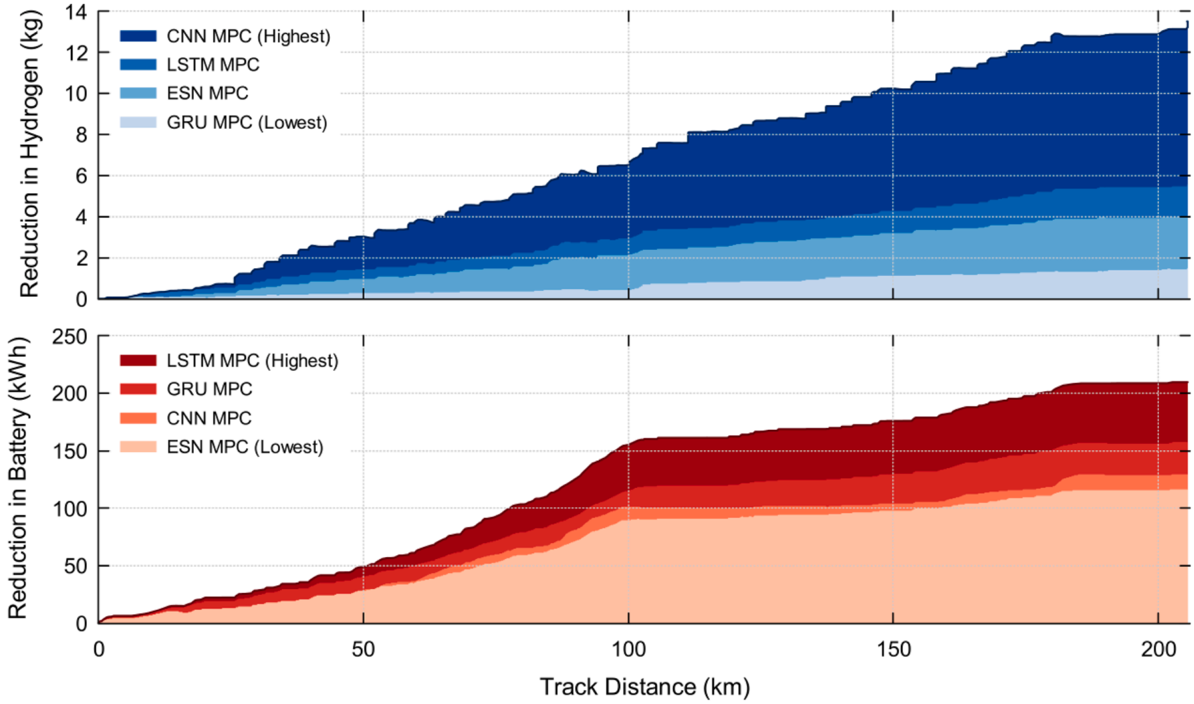

Fig. 10. Cumulative reduction in (top) hydrogen and (bottom) battery-energy consumption along the Brescia-Iseo-Edolo round trip, for each learning-augmented MPC relative to the baseline MPC.

aggressively optimizing hydrogen consumption, indicating the fundamental trade-off between fuel economy and component life [6,61].

7.2.3. Battery stress

LSTM-MPC achieved the lowest battery stress at 202.0 MW, corresponding to a 15.3% reduction relative to the baseline of 238.5 MW. This improvement could be attributed to a smoother battery power profile enabled by the LSTM predictions, reducing cumulative cycling and aggressive fluctuations and thereby mitigating capacity reduction and impedance gains [8,9]. GRU-MPC and ESN-MPC yield similar battery stress levels of 208.0 MW (12.8% reduction), indicating moderate battery protection performance. CNN-MPC recorded 214.0 MW (10.3% reduction), outperforming fuel cell stress performance but still falling short of LSTM.

Consistently observed battery stress reduction across all learning-augmented strategies reflects an improvement in the adjustment of the power contribution of fuel cells and batteries, made possible by future demand forecasting and subsequent power distribution optimization. The change in battery state of charge throughout the round trip at the Brescia-Iseo-Edolo line is shown in Fig. 11.

7.2.4. Battery charge sustaining

All learning-augmented strategies achieved perfect charge-sustaining ($\Delta SOC = 0.00\%$) compared with a 0.12% deviation from the reference MPC. This improvement demonstrates that precise power-demand predictions enabled MPC controllers to balance battery charging and discharging throughout the round-trip operation, preventing progressive SOC drift that could undermine operational reliability. The achievement of charge sustaining across all architectures suggests that this objective is relatively insensitive to differences in prediction accuracy, since only moderate improvements in prediction accuracy provide sufficient information for MPC to maintain SOC regulation [11,12].

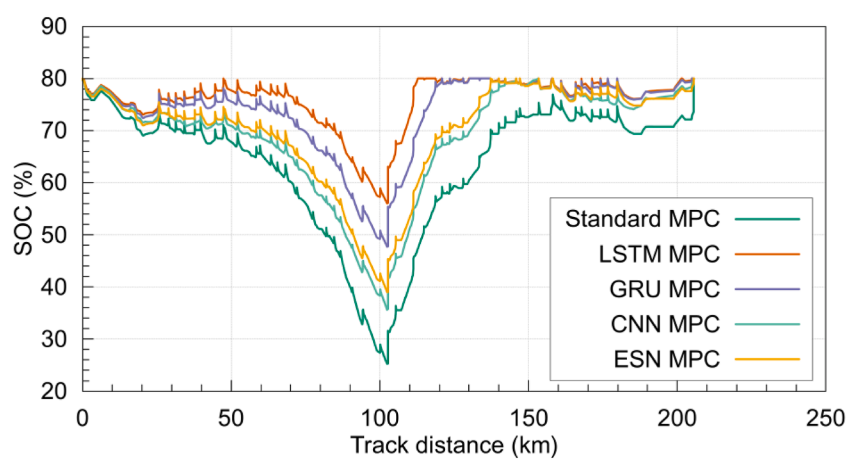

Fig. 11. Battery state-of-charge profile along the Brescia-Iseo-Edolo round trip for the baseline and learning-augmented MPC strategies.

7.2.5. Implications for the training line

The proposed EMS outputs in the training line demonstrate a fundamental inconsistency between operating objectives. CNN-MPC is ideal for applications that maximize hydrogen savings while prioritizing fuel economy over component life at the cost of increased fuel cell stress. LSTM-MPC is balanced, offering moderate hydrogen savings and maximum stress reduction, making it suitable for component-life-priority applications. GRU-MPC and ESN-MPC deliver comparable performance across all objectives. This trade-off reflects the conflicting nature of energy management goals. Aggressive fuel cell load following minimizes hydrogen consumption but increases stress, while conservative driving extends component life but sacrifices fuel economy. Therefore, the choice of strategy should depend on operational priorities, economic factors, and route characteristics.

7.3. Validation routes performance

Tables 3 and 4 present energy management performance on the two validation routes, assessing the generalization capabilities of the proposed learning-augmented MPC strategies. The corresponding variations in hydrogen and battery energy reductions, together with the SOC profiles along the routes, are plotted in Figs. 12 and 13.

7.3.1. Limone – Ventimiglia route analysis

On the Limone-Ventimiglia line, GRU-MPC achieved the highest hydrogen consumption reduction of 5.8%, demonstrating superior generalization compared to conservative performance (2.3%) on the training route. These improvements suggest that the pattern learned by GRU was particularly suited to the characteristics of the Limone-Ventimiglia route, including long distances between stops and predictable coastal terrain. Additionally, GRU-MPC recorded the lowest fuel cell stress at 24.5 MW, which shows effective component protection with a 5.8% reduction from baseline. However, it has a 30.2% SOC deviation, indicating the lowest charge sustaining performance among all strategies. This suggests GRU's optimization for hydrogen savings and stress reduction came at the expense of SOC balance on certain routes.

LSTM-MPC attained perfect charge sustaining (0.00% deviation) while maintaining moderate hydrogen savings (2.2%) and reduced battery stress. This balanced performance demonstrates that LSTM is particularly appropriate for operational deployments where charge sustaining is crucial. In addition, LSTM-MPC achieves the highest battery energy reduction of 37.8% on the Limone-Ventimiglia route. On this route, CNN-MPC performed poorly relative to the training line, with a hydrogen reduction rate of only 1.4%, and fuel cell stress increased to 30.5 MW (a 17.3% increase over baseline). This reduction indicates a lack of generalization, likely due to CNN-trained convolutional filters overfitting specific patterns in the training line and failing to generalize to the different topographic properties of Limone-Ventimiglia.

ESN-MPC obtained moderate performance (1.4% hydrogen reduction, 26.4% battery energy reduction, 7.3% SOC deviation) across all metrics. Relatively good charge-sustaining performance reflects ESN's short-term prediction accuracy, enabling effective SOC management even on new routes.

Table 3
Energy-management performance on the Limone-Ventimiglia route, comparing the standard MPC with the four learning-augmented strategies.

Strategy	H_2 Reduction (%)	Battery kWh Reduction (%)	FC Stress (MW)	Battery Stress (MW)	Δ SOC (%)
Standard MPC	Baseline (0)	Baseline (0)	26.0	175.6	23.2
LSTM-MPC	2.2	37.8	27.0	160.0	0.00
GRU-MPC	5.8	14.7	24.5	164.6	30.2
CNN-MPC	1.4	28.6	30.5	224.0	2.7
ESN-MPC	1.4	26.4	25.0	164.3	7.3

Table 4
Energy management performance on the Firenze-Faenza route, comparing the standard MPC with the four learning-augmented strategies.

Strategy	H_2 Reduction (%)	Battery kWh Reduction (%)	FC Stress (MW)	Battery Stress (MW)	Δ SOC (%)
Standard MPC	Baseline (0)	Baseline (0)	34.9	260.0	0.8
LSTM-MPC	7.5	3.8	36.3	232.2	0.07
GRU-MPC	13.4	6.7	29.4	240.0	18.4
CNN-MPC	0.0	21.5	53.7	331.0	15.1
ESN-MPC	5.3	34.4	31.4	240.9	0.01

7.3.2. Firenze - Faenza route analysis

GRU-MPC yielded a 13.4% reduction in hydrogen, recording the best fuel economy among all strategies on the Firenze-Faenza route, while LSTM-MPC achieved a 7.5% reduction. This improvement on the Firenze-Faenza route suggests that the long-term predictive capability of recurrent architecture is particularly effective on routes where power profile prediction is easier. GRU-MPC achieved a fuel cell stress of 29.4 MW, while LSTM-MPC resulted in 36.3 MW. Both models maintained low battery stress, with LSTM-MPC achieving 232.2 MW and GRU-MPC at 240.0 MW. However, GRU-MPC showed a SOC deviation of 18.4%, indicating that aggressive optimization for hydrogen reduction partially compromised SOC balance, whereas LSTM-MPC maintained a deviation of only 0.07%.

ESN-MPC reached the highest battery energy reduction, 34.4%, on the current route. These results demonstrate enhanced regenerative braking optimization capabilities, with precise short-term power management that responds quickly to braking events, maximizes energy recovery, and maintains a smooth charge/discharge profile. Notably, it also achieved near-perfect charge sustaining (0.01% SOC deviation), attributable to short-term prediction accuracy, enabling precise SOC management.

CNN-MPC failed on this route, with hydrogen reduction of 0.0%, while fuel cell stress deteriorated to 53.7 MW (53.9% increase) and battery stress to 331.0 MW (27.3% increase). This confirms CNN's limited ability to generalize to new routes and indicates that the learned convolutional patterns did not transfer effectively to the characteristics of the Firenze-Faenza route.

8. Discussion

8.1. Cross-route patterns and overall insights

Comparing the performance between routes reveals the generalization characteristics of each architecture. LSTM shows consistent intermediate performance on all routes, best stress reduction on the training route, and higher SOC stability on the Firenze-Faenza route. GRU demonstrated variable performance, outstanding on the Limone-Ventimiglia (5.8% hydrogen reduction) and Firenze-Faenza (13.4% hydrogen reduction), but underperforming on the training route (2.3% hydrogen reduction). CNN shows strong performance (20.3% hydrogen reduction) on the training route but fails on the validation routes, indicating severe overfitting. ESN shows moderate but consistent performance across all routes, with excellent charge-sustaining performance.

These patterns demonstrate that recurrent architectures (LSTM, GRU) generalize better to new routes than CNNs. The decline in tracking accuracy for the CNN across unseen routes stems from structural overfitting to the localized demand of the training line. Unlike sequential networks that leverage recurrent loops to model generalized temporal dynamics, a CNN uses fixed-size, localized kernel filters that extract

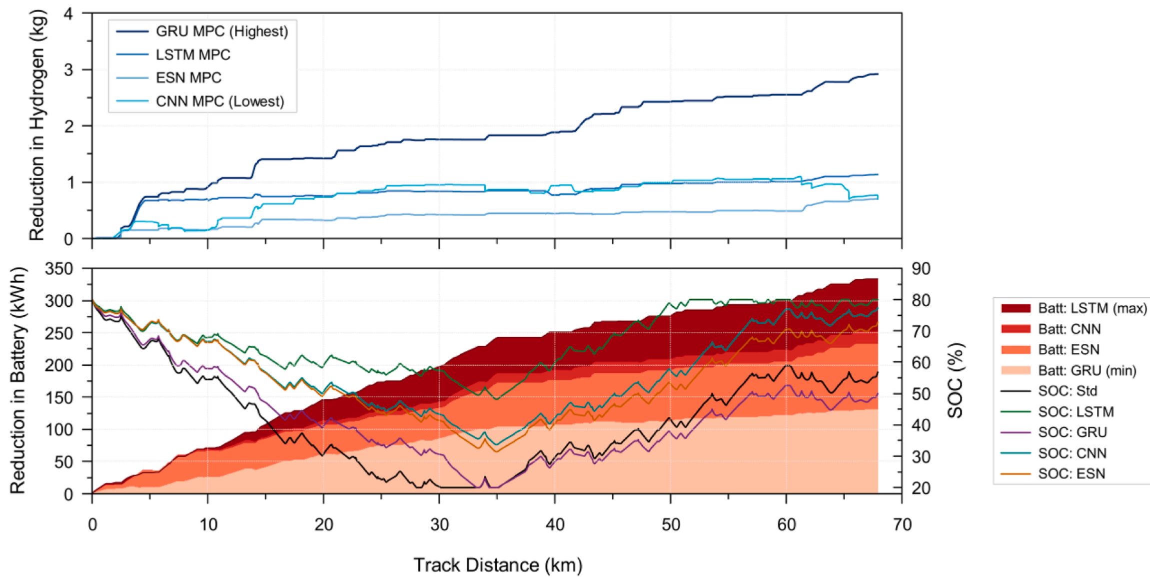

Fig. 12. EMS performance of the learning-augmented MPC strategies relative to standard MPC on the Limone-Ventimiglia route: hydrogen reduction (top) and battery-energy reduction with SOC profiles (bottom).

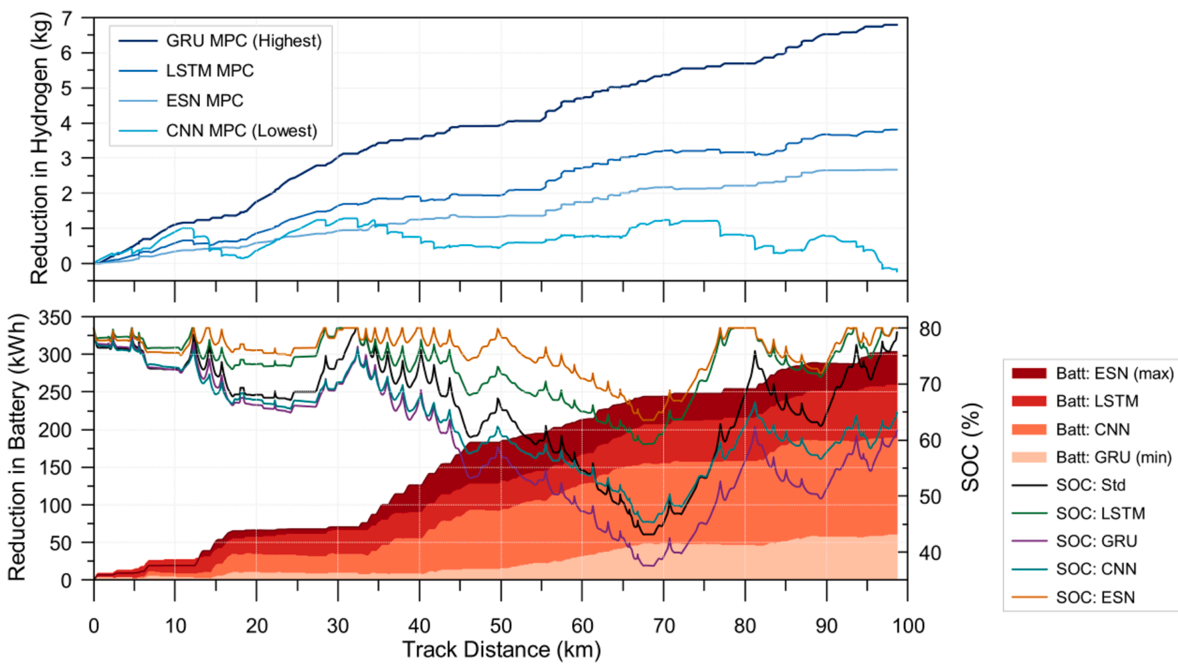

Fig. 13. EMS performance of the learning-augmented MPC strategies relative to standard MPC on the Firenze-Faenza route: hydrogen reduction (top) and battery-energy reduction with SOC profiles (bottom).

rigid, shape-specific feature configurations defined by the training route's historical cycles. When evaluated on unseen routes, these fixed-kernel mappings fail to generalize because the underlying localized demand frequencies, station spacing, and acceleration durations vary, leading to localized forecasting deviations that propagate into the optimization loop. Conversely, the ESN's consistent performance reflects the reservoir's general-purpose temporal processing capabilities, maintaining tracking stability across all profiles.

Furthermore, these cross-route findings reveal how forecasting errors interact with the closed-loop optimization strategy. Although multi-horizon prediction accuracy decays at later time steps, the receding-horizon mechanism executes only the initial control step, then refreshes the optimization space with real-time state feedback. This, combined with the high local trajectory stiffness provided by a heavily weighted state-of-charge objective, prevents forecasting errors from propagating or destabilizing the charge-sustaining performance across different routes.

To summarize, the energy-management analysis highlights four main findings. Low forecasting error alone does not guarantee optimal closed-loop model predictive control performance, underscoring the importance of alignment between predictor behavior and control objectives. Recurrent architectures, particularly Long Short-Term Memory and Gated Recurrent Unit, show stronger cross-route generalization in closed-loop energy-management performance, indicating better transfer of control-relevant temporal dynamics. Short-horizon predictive fidelity is especially important for maintaining charge-sustaining behavior. In contrast, the One-Dimensional Convolutional Neural Network is limited in multi-route deployment because of overfitting to the training route. These findings provide practical guidance for model selection. Fig. 14 summarizes trade-offs in a decision flowchart that maps operational priorities to the most appropriate predictor architecture.

Echo State Network is suitable when predictive precision and state-of-charge stability are prioritized. Long Short-Term Memory and Gated Recurrent Unit are more fitting when fuel economy, component protection, and cross-route generalization are the main priorities. The One-Dimensional Convolutional Neural Network should be used with caution in heterogeneous operating scenarios.

8.2. Computational cost and deployment feasibility

Fig. 15 compares the training time, memory usage, and parameter counts of the evaluated machine learning models. The reservoir model has the smallest hardware footprint, while the convolutional network requires the most computation. All four architectures can run on embedded processors used in Train Control and Management Systems (TCMS). Modern industrial TCMS units typically allocate several megabytes of non-volatile memory for diagnostics and prediction tasks. The maximum memory footprint of these prediction models represents a small overhead, preserving the vehicle's core processing bandwidth. Integrating the predictor is direct: it receives the current power demand and provides a 10-step forecast to the MPC solver at each cycle.

To ensure stable, deterministic behavior during rapid sampling, each network must operate within TCMS constraints. Gated recurrent networks require consistent onboard memory: LSTM uses up to 115 kB, while GRU reduces this to about 90 kB by merging internal states and eliminating the output gate. CNNs avoid temporal feedback through parallel multiply-accumulate steps but require up to 233 kB of memory and overfit spatially with 59,738 parameters, making fast training for new routes impractical. ESNs minimize the parameter count to 10,010 with an 80 kB cache by fixing the reservoir and training only the linear output layer. This lowers latency, enables online adaptation, and reduces retraining time to 6.3 minutes, compared with 15-18 minutes for gated models. All rely on receding-horizon optimization, executing only the initial action each cycle, which limits multi-step error and onboard latency.

A further practical consideration is software safety. The learning-augmented model predictive control framework retains a deterministic numerical solver as the decision layer; the deep learning component provides an informational forecast rather than a direct control action. This structural separation simplifies the safety case and supports compliance with the existing railway software safety integrity standard.

8.3. Limitations and future direction

The limitations in the present work can be identified as follows:

1. The deep learning models rely on deterministic predictions. Point predictions may not capture uncertainties, which can be useful for robust MPC design. Probabilistic deep learning predictions could enable optimization that accounts for uncertainty.
2. A fixed prediction horizon is adopted. Although the 10-step horizon aligns with common MPC practices, it may not be optimal for all routes or objectives. Adaptive horizon selection based on prediction confidence could further improve performance.
3. A simplified performance degradation model is used for component health analysis. The adopted stress metric provides a first-order approximation of component aging but does not capture complex phenomena such as temperature effects, SOC-dependent aging, and

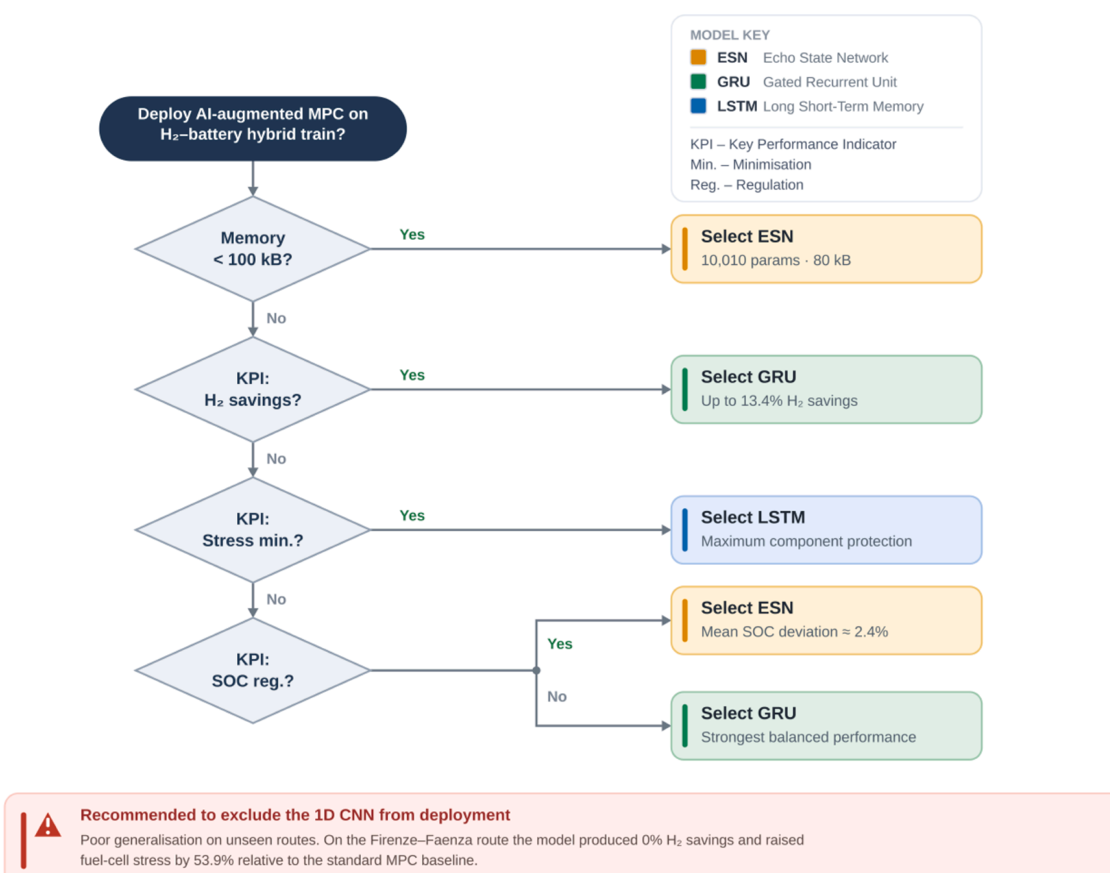

Fig. 14. Priority-based decision flowchart for predictor selection.

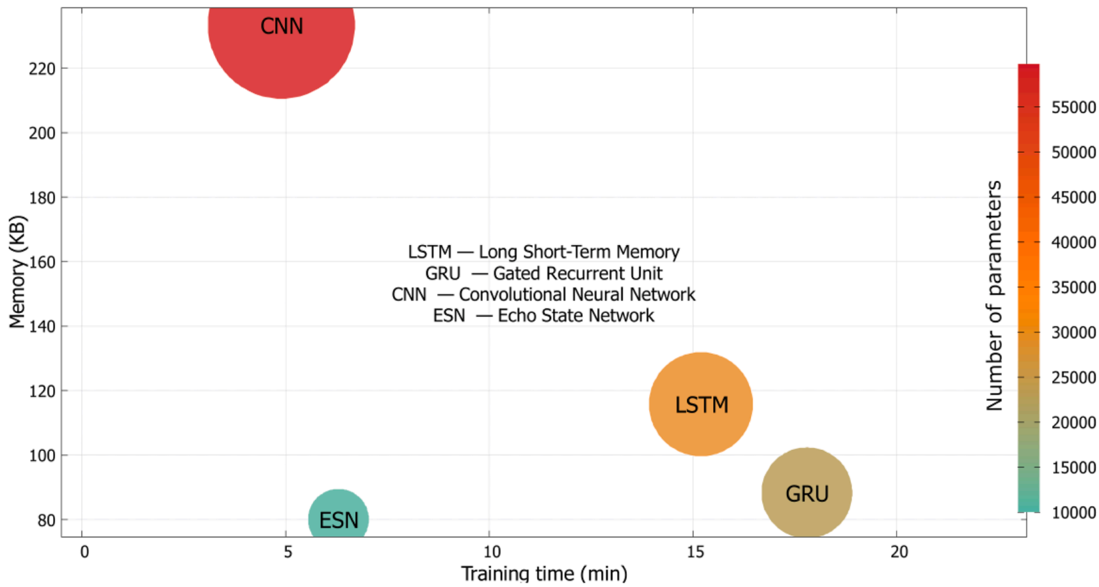

Fig. 15. Computational efficiency of the four models, comparing training time (x-axis), memory (y-axis), and parameter count (bubble size and color).

natural degradation with time. Integration with detailed electrochemical models could enable more accurate lifetime prediction.

And some of the potential extensions from this research are explained below:

1. Testing with hardware-in-the-loop to verify the performance of the proposed system under real-world conditions with sensor noise, communication delays, and safety constraints.
2. Incorporating the hydrogen fuel cost term and maintenance costs into the MPC objective functions to optimize operational economics over purely technical indicators.
3. Developing a hybrid architecture combining the short-term accuracy of ESN with the long-term capabilities of LSTM in a hierarchical structure to achieve better prediction and EMS performance.

9. Conclusion

This study examines the influence of demand-predictor architectures (Long Short-Term Memory, Gated Recurrent Units, One-Dimensional Convolutional Neural Networks, and Echo State Networks) on a closed-loop, learning-augmented Model Predictive Control framework for power splitting in hydrogen-battery hybrid trains. Evaluating all four networks on one training and two validation routes against a standard Model Predictive Control baseline yields three key insights:

1. Predictor architecture critically affects energy management, impacting hydrogen consumption, component stress, and State of Charge regulation independently of forecast accuracy. Notably, the most precise predictor (the Echo State Network) does not outperform others; gated recurrent architectures deliver superior closed-loop control. Selection should be based on closed-loop metrics over forecasting precision.
2. Systematic trade-offs enable tailored deployment: Gated Recurrent Units serve as a default baseline, balancing hydrogen savings and powertrain component health; Echo State Networks offer resource-efficient, near-optimal State of Charge regulation; Long Short-Term Memory excels in component stress minimization; and One-Dimensional Convolutional Neural Networks are not recommended due to poor cross-route generalization.
3. Out-of-distribution robustness is essential. While the One-Dimensional Convolutional Neural Network fails on unseen routes, the gated recurrent and Echo State network configurations remain robustly reliable. Evaluation must emphasize cross-route generalization over localized in-distribution fit.

These findings provide a rigorous, priority-driven framework that elevates demand-predictor selection from isolated algorithm benchmarking to a definitive deployment strategy to optimize the fuel cell-battery power split within an advanced, learning-augmented Model Predictive Control architecture for hydrogen-powered rail systems.

CRediT authorship contribution statement

Yosef Berhan Jember: Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation. **Massimo Santarelli:** Writing – review & editing, Supervision, Resources, Funding acquisition, Conceptualization. **Giuseppe Santangelo:** Writing – review & editing, Writing – original draft, Project administration, Funding acquisition. **Clemente Mirko:** Writing – review & editing, Writing – original draft, Supervision, Software, Investigation.

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 thank ALSTOM for providing information on the hydrogen train and the corresponding railway line, as well as for collaborative consultation with their experts. In cases where the company's input was inadequate, the authors used scientific methods and relevant literature to address gaps. This research was funded by the European Union, Next Generation EU - IPCEI Hydrogen 1 Hy2tech, Project ID: IPCEI-I1_0000003 CUP: B19J23000490004, coordinated by Alstom Ferroviaria S.p.A.

Data availability

The data that has been used is confidential.

References

[1] European Union Agency for Railways. Report on railway safety and interoperability in the EU 2024. LU: Publications Office; 2024. https://doi.org/10.2821/64343.

[2] Hoffrichter A, Hillmansen S, Roberts C. Conceptual propulsion system design for a hydrogen-powered regional train. IET Electr Syst Transp 2016;6:56–66. https://doi.org/10.1049/iet-est.2014.0049.

[3] Parola N, Peyrani G, Marocco P, Gandiglio M, Santarelli M. Impact of control strategies on the degradation of hybrid hydrogen-battery powertrains in railway applications. Int J Hydrog Energy 2025;177:151529. https://doi.org/10.1016/j.ijhydene.2025.151529.

[4] Ferrovie.it - Idrogeno, ANSFISA autorizza gli HMU.214 di Alstom per la Brescia-Edolo. Ferrovie.it n.d. https://www.ferrovie.it/portale/articoli/18019 (accessed June 21, 2026).

[5] Peng H, Li J, Löwenstein L, Hameyer K. A scalable, causal, adaptive energy management strategy based on optimal control theory for a fuel cell hybrid railway vehicle. Appl Energy 2020;267:114987. https://doi.org/10.1016/j.apenergy.2020.114987.

[6] Ferrara A, Jakubek S, Hametner C. Energy management of heavy-duty fuel cell vehicles in real-world driving scenarios: robust design of strategies to maximize the hydrogen economy and system lifetime. Energy Convers Manag 2021;232:113795. https://doi.org/10.1016/j.enconman.2020.113795.

[7] Wu J, Yuan XZ, Martin JJ, Wang H, Zhang J, Shen J, et al. A review of PEM fuel cell durability: degradation mechanisms and mitigation strategies. J Power Sources 2008;184:104–19. https://doi.org/10.1016/j.jpowsour.2008.06.006.

[8] Xu B, Oudalov A, Ulbig A, Andersson G, Kirschen DS. Modeling of lithium-ion battery degradation for cell life assessment. IEEE Trans Smart Grid 2018;9:1131–40. https://doi.org/10.1109/TSG.2016.2578950.

[9] Barré A, Deguilhem B, Grolleau S, Gérard M, Suard F, Riu D. A review on lithium-ion battery ageing mechanisms and estimations for automotive applications. J Power Sources 2013;241:680–9. https://doi.org/10.1016/j.jpowsour.2013.05.040.

[10] D'Ovidio G, Ometto A, Valentini O. A novel predictive power flow control strategy for hydrogen city rail train. Int J Hydrog Energy 2020;45:4922–31. https://doi.org/10.1016/j.ijhydene.2019.12.067.

[11] Li Q, Liu P, Meng X, Zhang G, Ai Y, Chen W. Model prediction control-based energy management combining self-trending prediction and subset-searching algorithm for hydrogen electric multiple unit train. IEEE Trans Transp Electrification 2022;8:2249–60. https://doi.org/10.1109/TTE.2022.3149479.

[12] Xie S, Hu X, Xin Z, Brighton J. Pontryagin's Minimum principle based model predictive control of energy management for a plug-in hybrid electric bus. Appl Energy 2019;236:893–905. https://doi.org/10.1016/j.apenergy.2018.12.032.

[13] Peng J, He H, Xiong R. Rule based energy management strategy for a series–parallel plug-in hybrid electric bus optimized by dynamic programming. Appl Energy 2017;185:1633–43. https://doi.org/10.1016/j.apenergy.2015.12.031.

[14] Sundstrom O, Guzzella L. A generic dynamic programming Matlab function. In: 2009 IEEE Int. Conf. Control Appl. St. Petersburg, Russia: IEEE; 2009. p. 1625–30. https://doi.org/10.1109/CCA.2009.5281131.

[15] Li X, Wang Y, Yang D, Chen Z. Adaptive energy management strategy for fuel cell/battery hybrid vehicles using Pontryagin's minimal Principle. J Power Sources 2019;440:227105. https://doi.org/10.1016/j.jpowsour.2019.227105.

[16] Kim Namwook, Cha Sukwon, Peng Huei. Optimal control of hybrid electric vehicles based on Pontryagin's minimum principle. IEEE Trans Control Syst Technol 2011;19:1279–87. https://doi.org/10.1109/TCST.2010.2061232.

[17] Musardo C, Rizzoni G, Guezennec Y, Staccia B. A-ECMS: an adaptive algorithm for hybrid electric vehicle energy management. Eur J Control 2005;11:509–24. https://doi.org/10.3166/ejc.11.509-524.

[18] Tian X, He R, Sun X, Cai Y, Xu Y. An ANFIS-based ECMS for energy optimization of parallel hybrid electric bus. IEEE Trans Veh Technol 2020;69:1473–83. https://doi.org/10.1109/TVT.2019.2960593.

[19] Oh Y, Ryu J, Kim J, Lee H. A study on energy optimization strategy for hydrogen fuel cell train using ECMS with model predictive control. J Electr Eng Technol 2025;20:5667–78. https://doi.org/10.1007/s42835-025-02419-8.

[20] Carignano M, Roda V, Costa-Castello R, Valino L, Lozano A, Barreras F. Assessment of energy management in a fuel cell/battery hybrid vehicle. IEEE Access 2019;7:16110–22. https://doi.org/10.1109/ACCESS.2018.2889738.

[21] Rawlings JB, Mayne DQ, Diehl M. Model predictive control: theory, computation, and design. 2nd edition. Madison, Wisconsin: Nob Hill Publishing; 2017.

[22] Deng K, Fang T, Feng H, Peng H, Löwenstein L, Hameyer K. Hierarchical eco-driving and energy management control for hydrogen powered hybrid trains. Energy Convers Manag 2022;264:115735. https://doi.org/10.1016/j.enconman.2022.115735.

[23] Zhou Y, Ravey A, Péra M-C. A survey on driving prediction techniques for predictive energy management of plug-in hybrid electric vehicles. J Power Sources 2019;412:480–95. https://doi.org/10.1016/j.jpowsour.2018.11.085.

[24] Chatterjee D, Biswas PK, Sain C, Roy A, Ahmad F, Rahul J. Bi-LSTM predictive control-based efficient energy management system for a fuel cell hybrid electric vehicle. Sustain Energy Grids Netw 2024;38:101348. https://doi.org/10.1016/j.segan.2024.101348.

[25] Li Q, Xun Q, Yang H. A learning model predictive control framework with long short-term memory for energy management in fuel cell hybrid electric buses. IEEE Trans Ind Appl 2026;62:3046–59. https://doi.org/10.1109/TIA.2025.3618223.

[26] Xun Q, Li Q, Yang HA. Learning model predictive controller for energy management in fuel cell hybrid electric vehicles. In: IEEE Transp. Electrification Conf. Expo ITEC. Chicago, IL, USA: IEEE; 2024. p. 1–6. https://doi.org/10.1109/ITEC60657.2024.10599030.

[27] Cui J, Hou S, Zhu Z, Liu Y, Gao J. Energy management strategy of fuel cell electric vehicles using learning-based EMPC. IEEE Trans Ind Appl 2026:1–11. https://doi.org/10.1109/TIA.2026.3656132.

[28] Jia C, He H, Zhou J, Li J, Wei Z, Li K. Learning-based model predictive energy management for fuel cell hybrid electric bus with health-aware control. Appl Energy 2024;355:122228. https://doi.org/10.1016/j.apenergy.2023.122228.

[29] Song K, Huang X, Xu H, Sun H, Chen Y, Huang D. Model predictive control energy management strategy integrating long short-term memory and dynamic programming for fuel cell vehicles. Int J Hydrog Energy 2024;56:1235–48. https://doi.org/10.1016/j.ijhydene.2023.12.245.

[30] Li F, Wang X, Bao X, Wang Z, Li R. Energy management strategy for fuel cell vehicles based on online driving condition recognition using dual-model predictive control. Sensors 2024;24:7647. https://doi.org/10.3390/s24237647.

[31] Jia C, Liu W, Chau KT, He H, Zhou J, Niu S. Passenger-aware reinforcement learning for efficient and robust energy management of fuel cell buses. eTransportation 2026;27:100537. https://doi.org/10.1016/j.etran.2025.100537.

[32] Jia C, Liu W, He H, Chau KT. Decoupled safety supervision empowering efficient and safe energy management for fuel cell vehicles. Npj Sustain Mobil Transp 2026;3:16. https://doi.org/10.1038/s44333-026-00087-3.

[33] Deng K, Peng H, Dirkes S, Gottschalk J, Ünlübayir C, Thul A, et al. An adaptive PMP-based model predictive energy management strategy for fuel cell hybrid railway vehicles. eTransportation 2021;7:100094. https://doi.org/10.1016/j.etran.2020.100094.

[34] Deng K, Liu Y, Hai D, Peng H, Löwenstein L, Pischinger S, et al. Deep reinforcement learning based energy management strategy of fuel cell hybrid railway vehicles considering fuel cell aging. Energy Convers Manag 2022;251:115030. https://doi.org/10.1016/j.enconman.2021.115030.

[35] Olmos J, Iparragirre U, Gaztañaga H. Comparative study of artificial intelligence based energy management strategies for hydrogen railway vehicles. In: 2024 IEEE Veh. Power Propuls. Conf. VPPC. Washington, DC, USA: IEEE; 2024. p. 1–8. https://doi.org/10.1109/VPPC63154.2024.10755440.

[36] Jaeger H, Haas H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science 2004;304:78–80. https://doi.org/10.1126/science.1091277.

[37] Lukoševičius M, Jaeger H. Reservoir computing approaches to recurrent neural network training. Comput Sci Rev 2009;3:127–49. https://doi.org/10.1016/j.cosrev.2009.03.005.

[38] Matzner F, Mráz F. Locally connected Echo State networks for time series forecasting. In: Thirteen Int Conf Learn Represent; 2025.

[39] Peyrani G, Marocco P, Gandiglio M, Cherchi P, Santarelli M. Techno-economic modeling framework to assess the feasibility of hydrogen-powered trains on non-electrified routes. J Power Sources 2025;652:237677. https://doi.org/10.1016/j.jpowsour.2025.237677.

[40] Lukaszewicz P. Energy consumption and running time for trains: modelling of running resistance and driver behaviour based on full scale testing n.d.

[41] Davis WJ. The tractive resistance of electric locomotives and cars. 1926.

[42] Tan C, Stecca M, Soeiro TB, Dong J, Bauer P. Performance evaluation of an electric vehicle traction drive using Si/SiC hybrid switches. In: 2021 IEEE 19th Int. Power Electron. Motion Control Conf. PEMC. Gliwice, Poland: IEEE; 2021. p. 278–83. https://doi.org/10.1109/PEMC48073.2021.9432574.

[43] Chakravarthy BK, Power Sree Lakshmi G. Savings with all SiC inverter in electric traction applications. In: E3S Web Conf. 87; 2019, 01014. https://doi.org/10.1051/e3sconf/20198701014.

[44] Larminie J, Dicks AL. Fuel cell systems explained. 2. Aufl., Nachdruck. Chichester Weinheim: Wiley; 2009.

[45] O'Hayre RP, Cha S-W, Colella WG, Prinz FB, editors. Fuel cell fundamentals. 3rd ed. Hoboken, New Jersey: John Wiley & Sons Inc; 2016.

[46] Kandidayeni M, Macias A, Khalatbarisoltani A, Boulon L, Kelouwani S. Benchmark of proton exchange membrane fuel cell parameters extraction with metaheuristic optimization algorithms. Energy 2019;183:912–25. https://doi.org/10.1016/j.energy.2019.06.152.

[47] Pukrushpan JT, Stefanopoulou AG, Peng H. Control of fuel cell power systems: principles, modeling, analysis and feedback design. 1st ed. London: Springer; 2004.

[48] Borup R, Meyers J, Pivovar B, Kim YS, Mukundan R, Garland N, et al. Scientific aspects of polymer electrolyte fuel cell durability and degradation. Chem Rev 2007;107:3904–51. https://doi.org/10.1021/cr050182l.

[49] Jouin M, Gouriveau R, Hissel D, Péra M-C, Zerhouni N. Degradations analysis and aging modeling for health assessment and prognostics of PEMFC. Reliab Eng Syst Saf 2016;148:78–95. https://doi.org/10.1016/j.ress.2015.12.003.

[50] De Bruijn FA, Dam VAT, Janssen GJM. Review: durability and degradation issues of PEM fuel cell components. Fuel Cells 2008;8:3–22. https://doi.org/10.1002/fuce.200700053.

[51] Vetter J, Novák P, Wagner MR, Veit C, Möller K-C, Besenhard JO, et al. Ageing mechanisms in lithium-ion batteries. J Power Sources 2005;147:269–81. https://doi.org/10.1016/j.jpowsour.2005.01.006.

[52] Wang J, Liu P, Hicks-Garner J, Sherman E, Soukiazian S, Verbrugge M, et al. Cycle-life model for graphite-LiFePO4 cells. J Power Sources 2011;196:3942–8. https://doi.org/10.1016/j.jpowsour.2010.11.134.

[53] Nguyen NP, Duong TA, Jan P. Strategies of multi-step-ahead forecasting for Chaotic Time Series using autoencoder and LSTM Neural networks: a comparative study. In: Proc. 2023 5th Int. Conf. Image Process. Mach. Vis. Macau China: ACM; 2023. p. 55–61. https://doi.org/10.1145/3582177.3582187.

[54] Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput 1997;9:1735–80. https://doi.org/10.1162/neco.1997.9.8.1735.

[55] LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015;521:436–44. https://doi.org/10.1038/nature14539.

[56] Cho K, van Merrienboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, et al. Learning phrase representations using RNN encoder-decoder for statistical machine translation 2014. https://doi.org/10.48550/ARXIV.1406.1078.

[57] Chung J, Gulcehre C, Cho K, Bengio Y. Empirical evaluation of gated recurrent neural networks on sequence modeling 2014. https://doi.org/10.48550/ARXIV.1412.3555.

[58] Bai S, Kolter JZ, Koltun V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling 2018. https://doi.org/10.48550/arXiv.1803.01271.

[59] Kiranyaz S, Avci O, Abdeljaber O, Ince T, Gabbouj M, Inman DJ. 1D convolutional neural networks and applications: a survey. Mech Syst Signal Process 2021;151:107398. https://doi.org/10.1016/j.ymssp.2020.107398.

[60] Gallicchio C, Micheli A, Pedrelli L. Deep reservoir computing: a critical experimental analysis. Neurocomputing 2017;268:87–99. https://doi.org/10.1016/j.neucom.2016.12.089.

[61] Fletcher T, Thring R, Watkinson M. An Energy Management strategy to concurrently optimise fuel consumption & PEM fuel cell lifetime in a hybrid vehicle. Int J Hydrog Energy 2016;41:21503–15. https://doi.org/10.1016/j.ijhydene.2016.08.157.