
Article

LSTM-Based Prediction of System Dynamics in a Commercial AEM Water Electrolyzer Under Target Hydrogen Production Rates

Kuan-Yu Hsu, and Jenn-Kun Kuo *

Department of Mechanical and Electro-Mechanical Engineering, National Sun Yat-sen University, Kaohsiung 804201, Taiwan

* Correspondence: jenn.kun@gmail.com; Tel.: +886-7-5252000 (ext. 4288); Fax: +886-7-5254299

Received: 19 October 2025; Revised: 27 January 2026; Accepted: 29 January 2026; Published: 3 February 2026

Abstract: Commercial anion exchange membrane water electrolyzers (AEMWE) typically do not provide detailed information about internal system parameters or stack characteristics, making physics-based modeling difficult and limiting their integration into customized digital simulation environments. Moreover, electrochemical and multiphysics models often require extensive computational resources, and parameter optimization can be prohibitively expensive. This study investigates the use of a long short-term memory (LSTM) network to predict the dynamic response of a commercial AEMWE under varying hydrogen production rates. The goal is to establish a data-driven model that replicates the operational logic of the real system, enabling the prediction of stack current, stack voltage, and hydrogen flow rate directly from the target production rate with significantly lower computational cost while preserving the transient behavior of the device. Model validation shows that the proposed LSTM achieves RMSE values of 0.48 for stack current, 0.92 for stack voltage, and 5.22 for hydrogen flow rate. The model successfully captures the overall dynamic trends and demonstrates that hydrogen flow rate and stack voltage can be inferred from the target production rate and stack current. Training is computationally efficient, allowing rapid model development. Although undershoot occurs during rapid decreases in production rate due to the smoothing characteristics of LSTM, the overall prediction error remains within acceptable bounds. The results highlight the potential of data-driven modeling for fast and practical AEMWE system representation, supporting future development of digital twins and accelerating applications in green hydrogen technologies.

Keywords: anion exchange membrane water electrolysis; long short-term memory; deep learning; digital twin; dynamic modeling

1. Introduction

Since the Industrial Revolution, the atmospheric concentration of carbon dioxide has increased by more than 50%, intensifying the greenhouse effect and causing a continuous rise in global surface temperatures. In response to the rapidly growing energy demand, humanity is actively seeking cleaner and more sustainable energy solutions. Hydrogen energy, with its high energy density and zero-carbon characteristics, is regarded as one of the most promising energy storage technologies. Hydrogen (H_2) can be transported through compression, liquefaction, or chemical storage, and also serves as an important medium for long-term and seasonal energy storage.

To achieve net-zero carbon emissions, global hydrogen demand is projected to exceed 500 million tons per year by 2050 [1]. Facilitating a sustainable energy transition by integrating electrolysis hydrogen production systems with small wind turbines and photovoltaic (PV) systems has emerged as an effective solution [2]. However, due to the intermittent nature of renewable energy, understanding the dynamic behavior of electrolyzers under various operating temperatures and power inputs is crucial for overall system efficiency. Current water-electrolysis-based hydrogen production technologies can be categorized by operating temperature into high-temperature Solid Oxide Electrolysis Cells (SOECs) and low-temperature Proton Exchange Membrane Electrolysis Cells (PEMECs), Alkaline Water Electrolysis (AWE), and Microbial Electrolysis Cells (MECs) [3].

Among these, AWE technology can be traced back to 1789, when Adriaan and colleagues first succeeded in decomposing water into hydrogen and oxygen using an electrostatic generator [4]. AWE is currently the most mature hydrogen production technology [5] and has already achieved commercial-scale deployment worldwide.



Copyright: © 2026 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Publisher's Note: Scilight stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

However, AWE suffers from several drawbacks, including low current density, slow reaction rates, lower gas purity, slow start-up, and corrosion or maintenance issues associated with the KOH electrolyte, all of which reduce overall energy efficiency [6]. To overcome the limitations of AWE, polymer membrane electrolysis technology was developed in the 1960s and applied in aerospace applications. This technology offers high current density, high efficiency, rapid response, compact system design, low operating temperature (20–80 °C), and the ability to produce ultra-pure hydrogen [7]. Although PEMECs provide high efficiency and high current density, their reliance on precious metal catalysts results in significantly higher costs compared to AWE. Conversely, AWE is more cost-effective but is constrained by liquid electrolytes and limited dynamic performance. To address the limitations of both technologies, AEMWE emerged as an alternative. AEMWE utilizes earth-abundant catalytic materials to reduce capital costs while maintaining high energy efficiency and strong scalability for industrial applications [8].

Developing a multiphysics model for an electrolyzer involves the coupling of heat transfer, mass transfer, and electrochemical reactions, resulting in high model complexity and substantial computational cost, making it unsuitable for dynamic simulations [9]. At the cell and stack scales, multiscale models constructed using computational fluid dynamics (CFD) require significant computational resources, limiting their applicability for fast simulation or real-time control [10]. Industrial-scale electrolyzer modeling further requires numerous internal parameters [11], many of which are not accessible in commercial systems, making such simulations difficult to apply in practice. Therefore, given the high computational cost of multiphysics and CFD-based models and the limited availability of internal parameters in commercial devices, machine-learning-based modeling has emerged as a highly promising alternative, particularly for dynamic prediction and digital-twin applications.

Reference [12] compared multiple machine-learning models—including Elman neural networks, back-propagation neural networks (BPNN), long short-term memory (LSTM), random forest (RF), support vector machine (SVM), bidirectional LSTM (bi-LSTM), Firefly-optimized Elman networks, and genetic-algorithm-optimized BPNN (GA-BP)—for PEMEC parameter prediction, demonstrating that machine learning can effectively capture nonlinear relationships among electrolyzer parameters. Reference [13] employed a convolutional neural network–long short-term memory (CNN-LSTM) model to predict the time-series evolution of electrolyzer voltage, achieving accurate degradation prediction across two datasets. Reference [14] utilized a NARX architecture based on feedforward neural networks (FNN) and LSTM to predict PEMEC performance, successfully capturing its nonlinear dynamic characteristics and confirming the capability of LSTM models for nonlinear dynamic prediction of electrolyzers.

While conventional statistical optimization and metaheuristic algorithms have achieved significant results in screening electrolyzer parameters [15], dynamic prediction of commercial AEM systems under rapid transient operation still requires deep learning architectures with strong temporal modeling capabilities to capture their nonlinear characteristics.

To overcome the limitations of traditional multiphysics models in terms of computational cost and parameter accessibility, this study aims to develop a data-driven model capable of accurately describing the dynamic behavior of a commercial AEMWE system. Using measurement data collected under various load conditions, an LSTM-based multivariable prediction framework is constructed to estimate key operating parameters, including current, voltage, and hydrogen flow rate. The model effectively captures the nonlinear and transient characteristics of AEMWE and demonstrates fast inference capability, making it a strong foundation for future digital-twin development and control-system design.

2. Methodology

2.1. System Description

The experiments in this study were conducted using a commercial Enapter AEM Electrolyzer (EL 2.1) (Figure 1), which is capable of producing up to 500 NL/h of hydrogen. The system operates with a nominal electrical consumption of approximately 2.4 kW, including a 2.2 kW electrolyzer stack and roughly 0.6 kW of thermal losses during steady operation. The detailed specifications are summarized in Table 1.

To operate the Enapter AEM system, the user sets the target hydrogen production rate through a dedicated smartphone application. After the command is issued, the system performs hydration and ramp-up procedures before reaching stable operation. During operation, key system variables—including stack voltage, stack current, and hydrogen flow rate—are automatically uploaded to a local web interface for monitoring and data logging.

Although variables such as stack voltage, stack current, and hydrogen flow rate exhibit approximately linear relationships under steady-state conditions, their behavior becomes highly dynamic during warm-up and transient

operation. These variables do not follow a fixed static relationship and therefore cannot be described by simple analytical formulas.

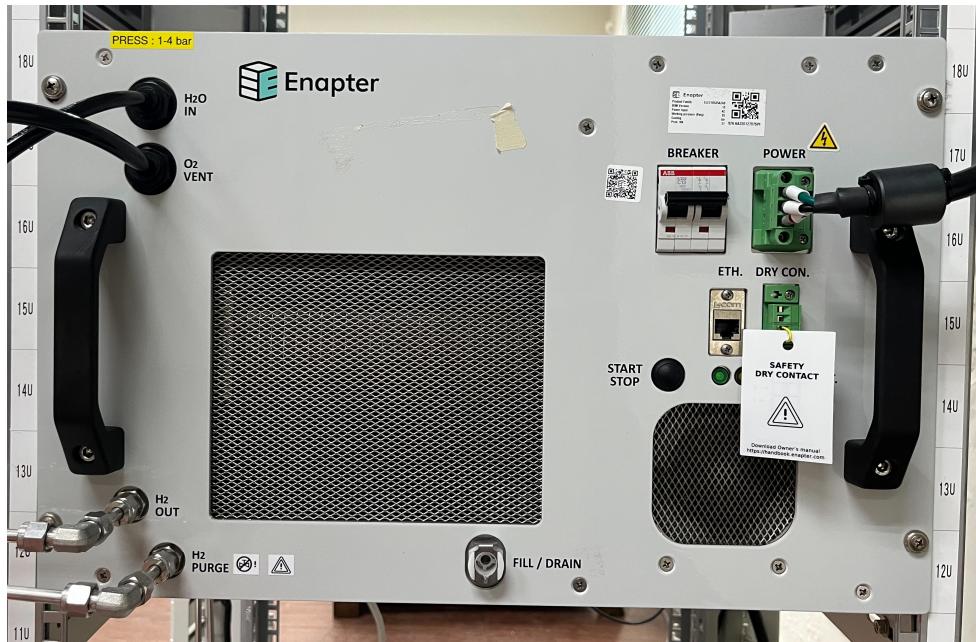


Figure 1. Enapter AEM electrolyzer EL2.1.

Table 1. Enapter AEM electrolyzer EL 2.1 specs.

Parameters	Value
Dimensions	W: 482 mm * D: 634 mm * H: 307 mm
Weight	55 kg
Production rate	500 NL/hr
Output pressure	Up to 35 barg
Hydrogen output purity	35 bar: ~99.9% (Impurities: ~1000 ppm H ₂ O) 8 bar: >1500 ppm H ₂ O
Nominal power consumption (beginning of life)	4.8 kWh/Nm ³
Operative power consumption	2400W
Power supply	200–240 VAC, 50/60 Hz
Maximum water input conductivity	20 μ S/cm at 25 °C
Water consumption	~400 mL/h
Water input pressure range	1–4 barg
Ambient operative temperature range	5–45 °C
Ambient operative humidity range	Up to 90% Rh, non-condensing
IP rating	IP 20
Control and monitoring	Fully automatic with Enapter's EMS, Modbus TCP via Ethernet

2.2. Data Preprocessing

In this study, the hydrogen production rate of the AEM electrolyzer was manually adjusted every 30 min with a step size of 10%, starting from 60%, increasing to 100%, and then decreasing back to 60%. The system recorded data every two seconds, resulting in a total of 14,165 samples over 275 min (Figure 2). From the system's automatically logged parameters, four variables—production rate, stack current, stack voltage, and hydrogen flow rate—were selected as the inputs and outputs for modeling the dynamic response of the electrolyzer.

The dataset was split chronologically, with 80% used for training and the remaining 20% for validation. In addition, an independent dynamic test was conducted under a different operating profile, where the production rate was adjusted every six minutes with the same 10% step size, cycling from 60% to 100% and back to 60% for two full cycles. This dataset was also sampled every two seconds, yielding 8798 samples over 170 min, and was used as an unseen test set to evaluate the model's generalization capability (Figure 3). Missing values in the dataset were handled using mean imputation. To mitigate the impact of scale differences among variables, all input and output features were standardized, which improves both training stability and prediction accuracy.

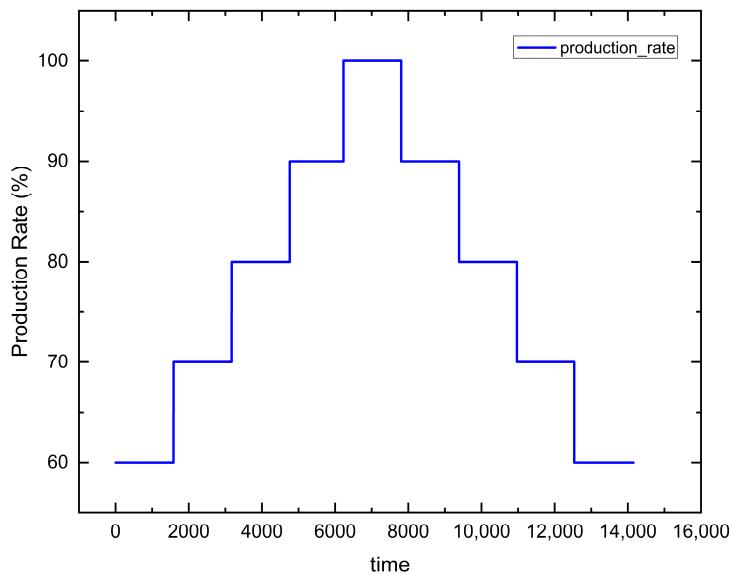


Figure 2. Training production rate cycle applied to AEM system.

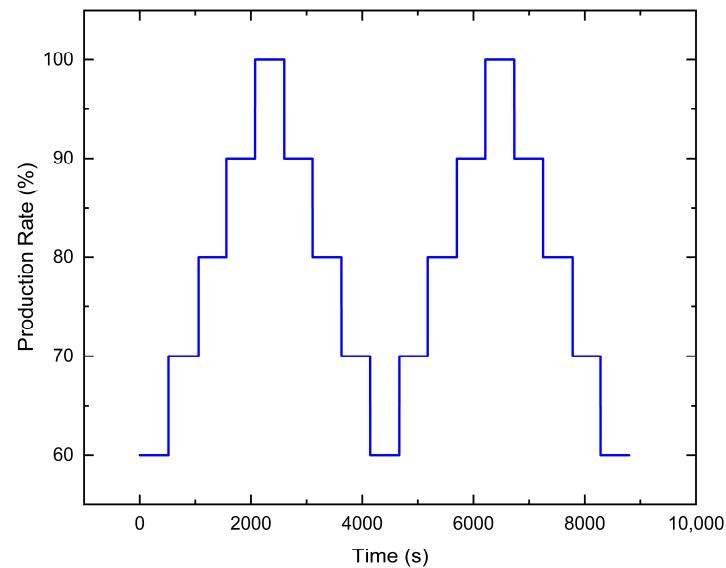


Figure 3. Test production rate cycle applied to AEM system for generalization evaluation.

To enable the LSTM model to learn the temporal dependencies of the system, a sliding-window approach was applied to transform the time-series data into supervised learning pairs. A window length of 30 time steps was used, where each sequence of 30 consecutive samples served as the model input to predict the system response at the next time step. This method preserves the sequential structure of the data and allows the model to capture the dynamic behavior of the AEM electrolyzer under varying production rate conditions.

2.3. LSTM Neural Network

Recurrent Neural Networks (RNNs) are capable of preserving temporal dependencies within sequential data, making them suitable for time-series modeling. However, conventional RNNs often suffer from the vanishing-gradient problem when processing long sequences. To address this limitation, the Long Short-Term Memory (LSTM) architecture introduces memory cells to replace the traditional hidden units, enabling the model to capture both short-term and long-term temporal dependencies (Figure 4).

For time series $X_T = (x_1, \dots, x_T)$, the LSTM receives an input x_t at each time step t and regulates information flow through three gating mechanisms: the input gate, forget gate, and output gate. The cell state C_t is updated based on the outputs of these gates, allowing the network to retain essential historical information throughout the sequence.

When a new input x_t enters the model, the forget gate determines which components of the previous hidden state h_{t-1} should be retained or discarded. Its output is computed as:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad (1)$$

where W_f and b_f denote the weight matrix and bias term of the forget gate.

The input gate controls whether the processed input x_t should be incorporated into the cell state:

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad (2)$$

where W_i and b_i represent the input gate's weights and bias.

The input signal is further transformed through a \tanh activation to generate a candidate information vector \tilde{C}_t , which is integrated into the updated cell state:

$$\tilde{C}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c) \quad (3)$$

where W_c and b_c correspond to the weights and bias associated with the candidate cell update.

The updated cell state C_t is computed as:

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (4)$$

Finally, the output gate determines the hidden state h_t by modulating the activated cell state:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (5)$$

$$h_t = o_t \cdot \tanh C_t \quad (6)$$

where W_o and b_o denote the output gate's weights and bias.

Each LSTM layer consists of N_c memory cells, and the number of cells determines the model's representational capacity. After processing the entire input sequence, the final hidden state is passed through a fully connected output layer to generate the system response at the next time step.

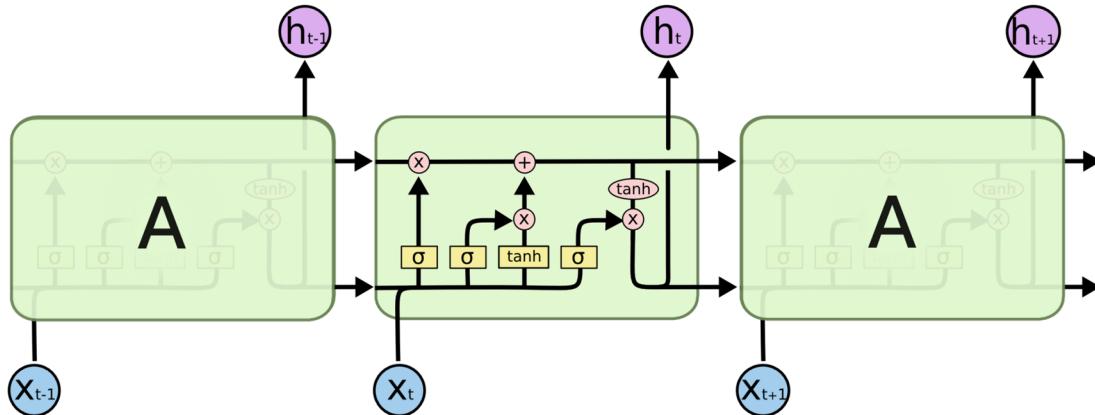


Figure 4. A chain of LSTM units [13].

2.4. Training Setup

The LSTM model in this study was implemented using the open-source PyTorch deep learning framework, and consists of a single hidden layer. The input data were transformed into fixed-length sequences through a sliding-window approach and trained using a mini-batch strategy. For each batch, the model performs a forward pass and computes the Mean Squared Error (MSE) as the loss function. Gradients across all time steps are then obtained through Backpropagation Through Time (BPTT), and the model parameters are updated using the Adaptive Moment Estimation (Adam) optimizer. This training procedure continues until the maximum number of epochs is reached or early stopping is triggered.

Regarding hyperparameter settings, the batch size was set to 32, the initial learning rate to 5×10^{-4} , the loss function to MSE, and the maximum number of training epochs to 100. To enhance training stability, the ReduceLROnPlateau scheduler was employed to adjust the global learning rate, with a decay factor of 0.5 and a patience of 5 epochs. When the validation loss plateaued, the scheduler automatically reduced the learning rate to facilitate more effective convergence. In addition, early stopping with a patience of 10 epochs was applied to

prevent overfitting and improve training efficiency by halting the process when the validation loss failed to improve for 10 consecutive epochs.

Optimizer selection plays a critical role in training RNN-based models. In this work, the Adam optimizer was adopted, which can be viewed as a combination of Momentum and RMSProp. Momentum accumulates the first-order moment of gradients, providing inertia that accelerates convergence and reduces oscillations, but may accumulate outdated gradient directions and lead to suboptimal updates. RMSProp adjusts the learning rate of each parameter based on the second-order moment of gradients, enabling more stable progress in parameter spaces with varying scales; however, the absence of a momentum term makes it prone to oscillation in narrow valleys or stagnation near saddle points. Adam integrates the advantages of both methods by incorporating momentum-based smoothing and adaptive per-parameter learning rates, thereby improving training stability and convergence efficiency.

At the end of each epoch, the model automatically saves the set of parameters that yields the lowest validation loss. After training, this best-performing model is used for inference on the test dataset to ensure optimal predictive performance.

3. Results and Discussion

3.1. Correlation Analysis and Feature Selection Justification

Based on the Pearson Correlation Coefficient (PCC) analysis (Figure 5), the hydrogen flow rate (*h2_flow*) and stack current (*i_stack*) exhibit a correlation coefficient of $r = 1.00$, indicating a perfectly linear relationship. In addition, the correlations between *h2_flow* and the target hydrogen production rate (*pro_rate*), as well as between *i_stack* and *pro_rate*, both reach $r = 0.99$, representing an almost perfectly linear relationship. In comparison, the correlation coefficients between the stack voltage (*v_stack*) and the aforementioned variables fall within the range of 0.96–0.97. Although slightly lower, these values still indicate a highly linear relationship.

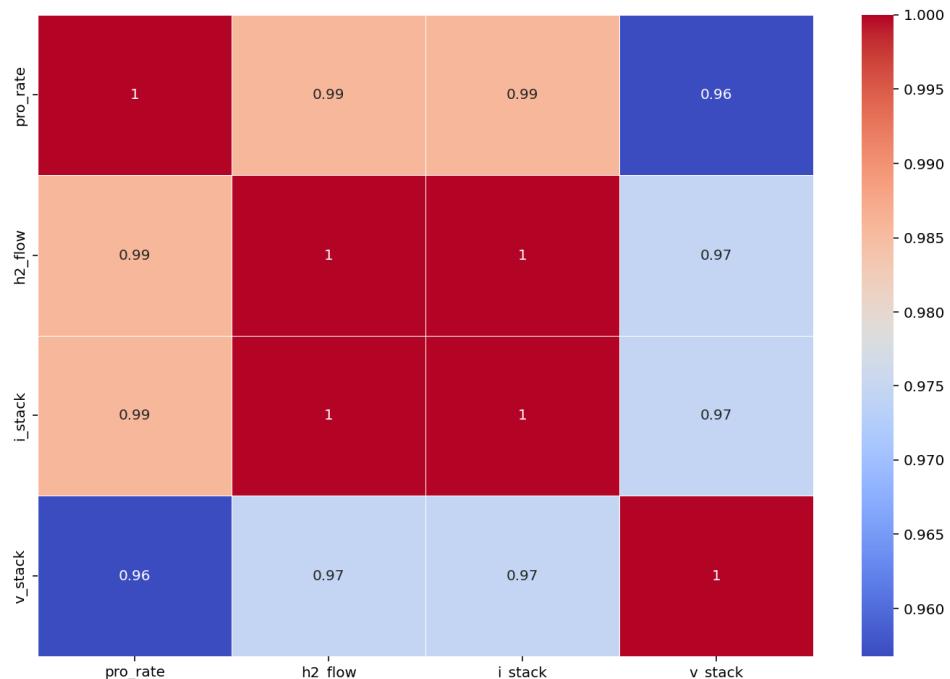


Figure 5. PCC correlation heatmap.

The perfectly linear relationship between *h2_flow* and *i_stack* is consistent with the reaction-rate behavior described by Faraday's law of electrolysis. According to Faraday's law:

$$n = \frac{Q}{zF} \quad (7)$$

where n is the amount of substance produced, Q is the total electric charge passing through the electrolyzer, z is the number of electrons transferred, and F is Faraday's constant. The total charge Q is given by the time integral of the current:

$$Q = \int I(t)dt \quad (8)$$

In this study, the data were recorded at a fixed sampling frequency, meaning that the time interval Δt is constant. Under this condition, the integral can be approximated as the summation of current values, implying a linear relationship between Q and I . Since the hydrogen production rate n is proportional to Q , the observed perfect linearity ($r = 1.00$) between h2_flow and i_stack is fully consistent with electrochemical theory.

Given the extremely high linear correlations among these variables, including all of them as model inputs would introduce substantial redundancy. Such multicollinearity can obscure the gradient direction during training, reduce parameter update stability, and ultimately degrade model robustness. Moreover, incorporating multiple highly correlated features unnecessarily increases model complexity and raises the risk of overfitting without providing additional predictive benefit. Therefore, this study selects the target hydrogen production rate (pro_rate) and stack current (i_stack) as the input features, representing the most informative and physically meaningful combination of operating variables.

3.2. Training Performance

The training loss curve of the LSTM model (Figure 6) shows that both the training loss and validation loss decrease rapidly during the initial phase, indicating that the model quickly learns the relationship between the hydrogen production rate and the stack current, stack voltage, and hydrogen flow rate. In the later stages, the losses continue to decrease smoothly, with only minor fluctuations that do not lead to divergence, demonstrating stable training behavior. Early stopping was triggered at epoch 41, preventing unnecessary over-training. Throughout the process, the validation loss remained slightly lower than the training loss, suggesting that the model did not exhibit overfitting.

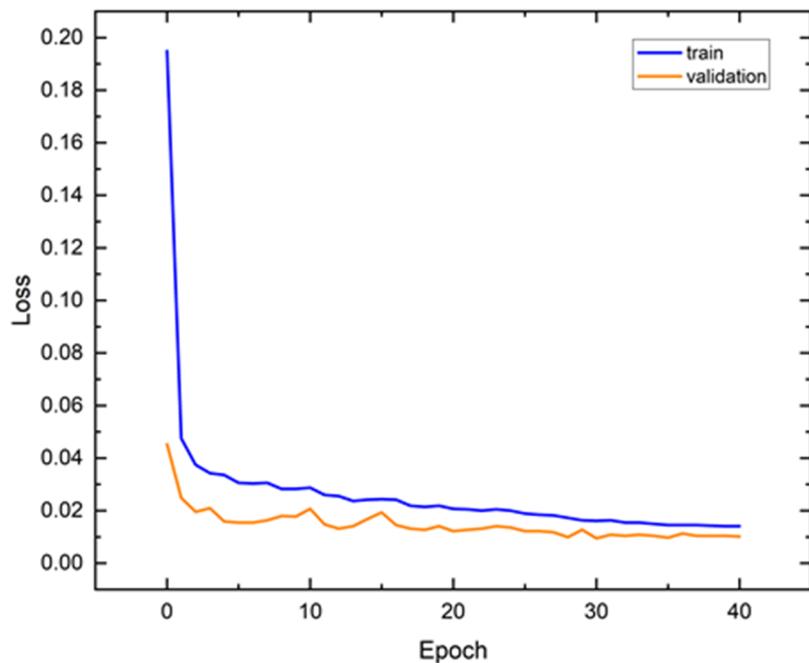


Figure 6. Loss curve.

During the training process, the loss curves exhibit smooth convergence, indicating that the model effectively captures the mapping between the input and output variables. This convergence behavior is consistent with the characteristics of the dataset: the input features (pro_rate and i_stack) and output variables exhibit strong linear relationships, enabling the model to rapidly grasp the dominant trends and achieve substantial loss reduction in the early epochs. Furthermore, the use of a relatively small initial learning rate ensures stable parameter updates; thus, the slight oscillations observed in later epochs do not compromise training stability or convergence.

With the early stopping mechanism in place, training is terminated once the validation loss ceases to improve, preventing overfitting while maintaining computational efficiency. Overall, the training curves demonstrate that the model achieves good convergence and generalization performance, providing a solid foundation for subsequent evaluation under dynamic operating conditions.

3.3. Test Results Under Dynamic Operating Conditions

Under dynamic operating conditions, the model demonstrates strong tracking performance across all three output variables. The predicted trajectories closely overlap with the true signals, and even the small oscillations embedded in the operating profile are accurately reproduced. In the region where the hydrogen production rate reaches 100%, the predicted hydrogen flow rate and stack current are slightly lower than the measured values, yet the overall trend remains consistent with the ground truth. No noticeable temporal lag is observed, indicating that the LSTM effectively follows rapid system transitions. (Figures 7–9).

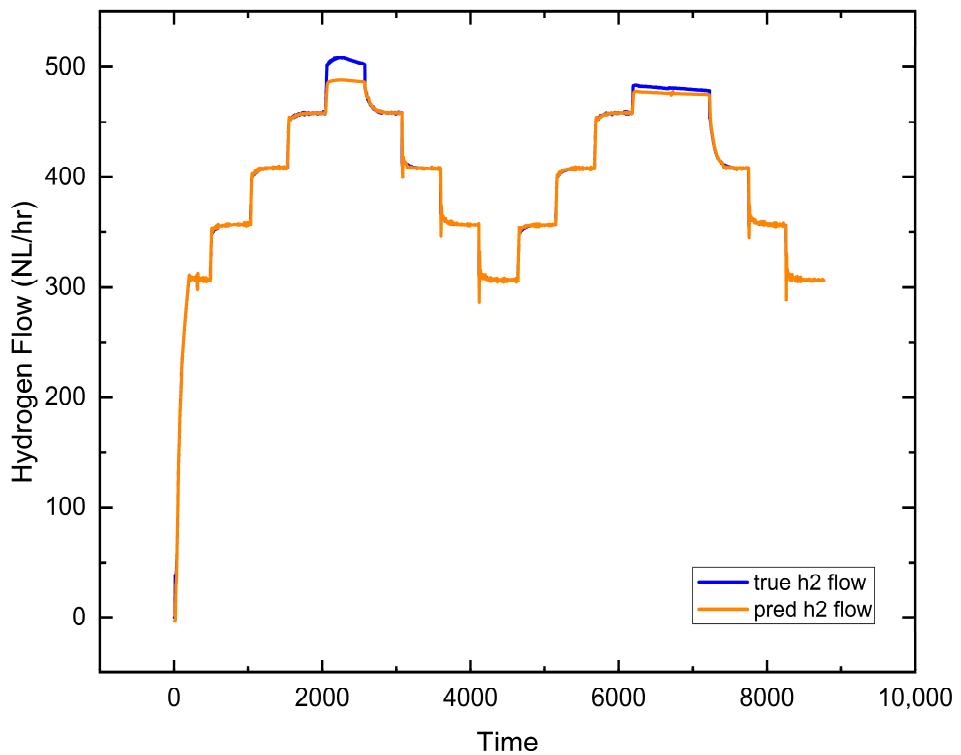


Figure 7. LSTM model prediction of hydrogen flow on unseen test data.

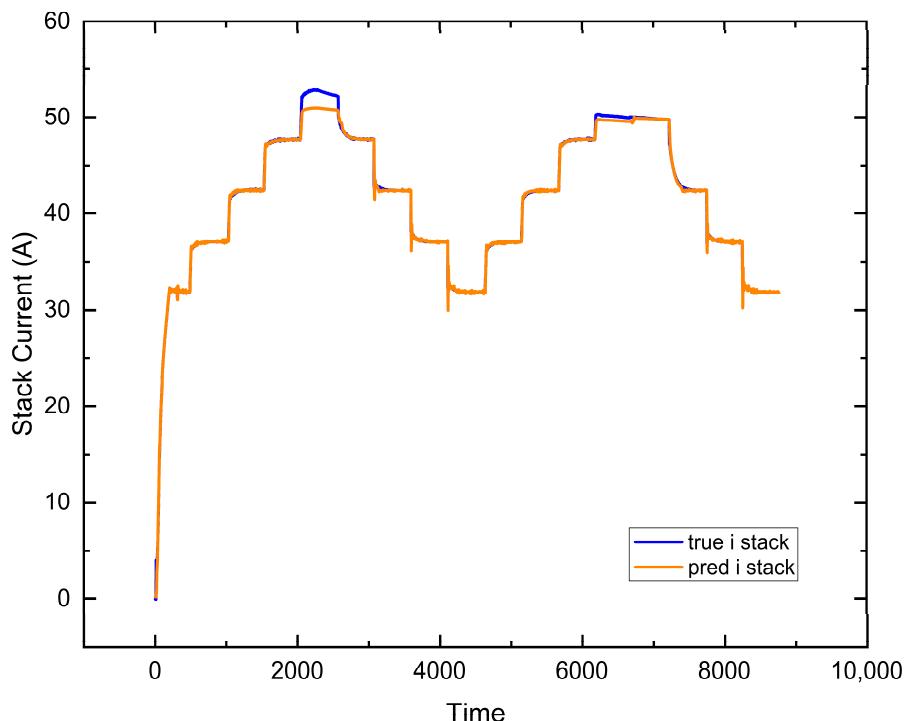


Figure 8. LSTM model prediction of stack current on unseen test data.

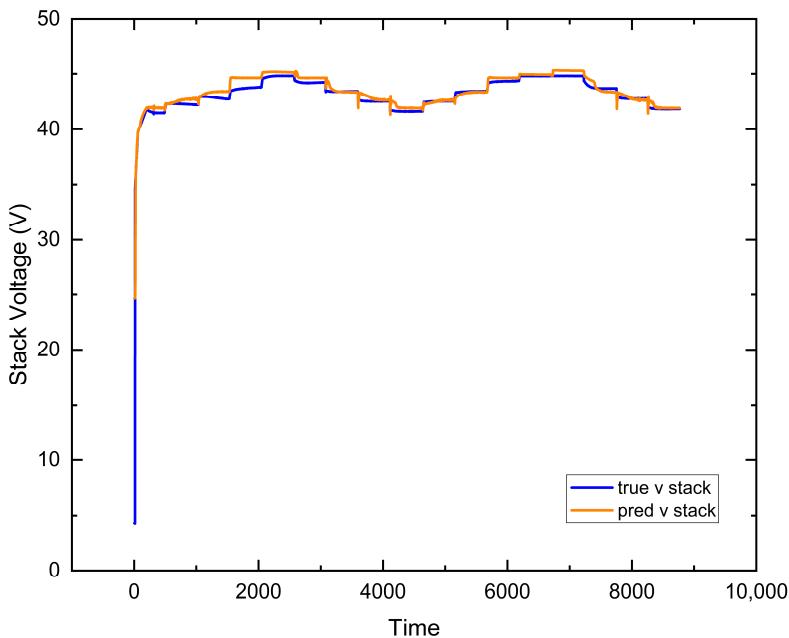


Figure 9. LSTM model prediction of stack voltage on unseen test data.

Across all variables, undershoot behavior commonly appears during sharp downward transitions, reflecting the model's conservative response when the system output drops abruptly. Hydrogen flow rate and stack current exhibit the most accurate tracking during normal operation, with deviations primarily concentrated in fast-decreasing segments. Although overshoot events are present, their magnitudes are generally small: the average overshoot is 0.66 for hydrogen flow (maximum 17.49, 1757 instances), 0.07 for stack current (maximum 1.65, 1463 instances), and 0.42 for stack voltage (maximum 20.39, 3283 instances). In contrast, undershoot is most pronounced in hydrogen flow, with an average deviation of -4.36 and a maximum of -41.97 (826 instances), while stack current and voltage exhibit milder undershoot patterns. Overall, the model successfully captures both the global trends and fine-scale dynamics of the system, exhibiting predictable overshoot and undershoot behaviors consistent with the rapid transitions in the operating conditions.

To quantitatively evaluate the model's prediction accuracy under dynamic conditions, RMSE, MAE, and MAPE were computed for all three output variables. Stack current exhibits the lowest prediction error, with an RMSE of 0.48 and an MAE of 0.21, indicating that the model accurately captures its dynamic behavior. Stack voltage shows slightly higher errors (RMSE 0.92, MAE 0.37), mainly due to nonlinear polarization effects and measurement noise, yet the overall prediction quality remains satisfactory.

In contrast, hydrogen flow rate yields significantly larger errors, with an RMSE of 5.22 and an MAE of 2.32, reflecting the greater difficulty in modeling its dynamic behavior. These error metrics align with the tracking analysis: stack current is the easiest variable to predict, followed by stack voltage, while hydrogen flow rate exhibits the largest errors due to its strong nonlinearity and rapid fluctuations.

From a physical perspective, the prediction difficulty varies across the three output variables. Stack current is a direct electrochemical output with clean measurements and fast response, enabling the model to reconstruct its dynamics with low error. Stack voltage is influenced by polarization behavior, humidity, temperature, and measurement noise, resulting in moderate nonlinearity and slightly higher errors, yet its overall tracking performance remains stable.

In contrast, although hydrogen flow rate is theoretically linear with respect to current, its dynamic behavior is affected by additional factors such as valve actuation delay, gas compressibility, manifold volume, pressure fluctuations, and low-flow measurement noise. These effects introduce strong nonlinearity and noise into the true signal. Consequently, even when current varies smoothly, hydrogen flow rate exhibits pronounced transient deviations during rapid increases and decreases, leading to significantly larger RMSE and MAE values.

Furthermore, the temporal smoothing nature of LSTM causes the model to produce undershoot during abrupt downward transitions, a phenomenon particularly evident in hydrogen flow due to its larger magnitude and faster rate of change. Overall, the differences in prediction error reflect the distinct physical dynamics and measurement characteristics of each variable rather than their theoretical linear relationships.

Overall, the combined analysis of tracking behavior, error metrics, and physical mechanisms demonstrates that the proposed model exhibits strong generalization capability under dynamic operating conditions. All three

output variables are accurately reconstructed, capturing both their major trends and transient variations without noticeable delay or instability. The low prediction errors in stack current and voltage indicate that the model effectively learns both linear and moderately nonlinear electrochemical behaviors. Although hydrogen flow rate shows larger errors due to its strong nonlinearity and measurement noise, the model still reproduces its rising and falling transitions with reasonable dynamic consistency.

In summary, the model maintains stable prediction quality across different operating regions and responds promptly to rapid input changes, demonstrating robust dynamic tracking performance and practical applicability. These results are consistent with the training-phase behavior, further confirming the model's reliability when applied to previously unseen dynamic data.

3.4. Discussion of Model Behavior

The model exhibits clear asymmetry between the rising and falling phases of the system dynamics. In the real system, the hydrogen flow rate decreases much more rapidly than it increases, and the LSTM requires the updated output information before it can fully adjust its prediction. As a result, the model tends to show delayed responses or overshoot during sharp downward transitions. This behavior reflects the inherent limitations of sequence-based models when handling abrupt changes, yet remains consistent with typical characteristics of time-dependent neural architectures.

The differences in prediction difficulty among the three output variables further highlight the distinct nature of their physical dynamics. Stack current is directly controlled by the production rate, making it the cleanest and most predictable signal. Stack voltage, in contrast, is inferred from the current, causing any current-related error to propagate and accumulate, in addition to the voltage's inherent sensitivity to polarization behavior and operating conditions. Hydrogen flow rate is the most challenging variable to predict, as its dynamics are influenced by control delays, measurement noise, pressure fluctuations, and gas compressibility, all of which introduce stronger nonlinearity and uncertainty into the signal.

Moreover, the structural characteristics of the LSTM also shape the model's dynamic behavior. Because the LSTM retains information from previous time steps and uses it to estimate future outputs, it is prone to delayed or overshooting responses when the system undergoes sudden and rapid changes. This tendency is particularly evident in variables with large magnitudes and fast downward transitions, such as the hydrogen flow rate.

3.5. Summary of Findings

The model demonstrates strong overall performance under dynamic operating conditions, accurately capturing the major trends of all three output variables and maintaining consistent behavior during rapid system transitions. Among the variables, stack current is predicted with the highest accuracy due to its direct dependence on the production rate, which provides a clean and well-defined control relationship. In contrast, hydrogen flow rate is the most challenging variable to predict, as its dynamics are influenced by control delays, measurement noise, and pressure-related effects. The most prominent model deviation occurs in the form of undershoot during downward transitions, reflecting the LSTM's characteristic response to abrupt decreases. Overall, the model is capable of rapidly reconstructing dynamic system behavior and shows strong potential for practical application.

4. Conclusions

This study developed a data-driven model capable of accurately predicting the dynamic behavior of water electrolyzers, demonstrating strong generalization performance on unseen test data. Without relying on complex physical formulations, the model effectively reconstructs key system dynamics and shows practical potential for dynamic operation and digital-twin applications. Although undershoot deviations remain during rapid downward transitions, the overall prediction quality remains stable.

The main contributions of this work are as follows:

- (1) A multivariable dynamic prediction model for water electrolyzers that captures nonlinear and transient behaviors effectively.
- (2) Demonstration that data-driven modeling can replace high-cost physical simulations while maintaining high accuracy with low computational demand.
- (3) Verification of strong generalization capability on previously unseen operating conditions.
- (4) A foundational framework for digital-twin and real-time monitoring applications in water electrolyzer systems.
- (5) Identification of LSTM limitations in fast-changing regimes, with future directions toward incorporating physical priors to enhance accuracy and interpretability.

Author Contributions: K.-Y.H.: conceptualization, methodology, software, data curation, writing—original draft preparation, visualization, investigation; J.-K.K.: supervision, validation, writing—reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Science and Technology Council of Taiwan under Contract No. NSTC 112-2923-E-110-002-MY3, NSTC 112-2221-E-017-MY3, NSTC-113-2923-E-008-007, NSTC 114-2622-E-110-007, NSTC-114-2923-E-110-1103-003-MY3.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to institutional restrictions and confidentiality agreements.

Conflicts of Interest: The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies: During the preparation of this work, the authors used Microsoft Copilot to assist with translation and English language polishing of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

1. Gevaert, S.; Pause, L.; Cezne, E.; et al. Green Hydrogen in the Global South: A Literature Review. Available online: https://www.researchgate.net/profile/Eric-Cezne/publication/366809109_Green_Hydrogen_in_the_Global_South_A_literature_review/links/63b2fb56097c7832ca84c636/Green-Hydrogen-in-the-Global-South-A-literature-review.pdf (accessed on 22 April 2025).
2. Huang, P.H.; Kuo, J.K.; Wu, Z.D. Applying Small Wind Turbines and a Photovoltaic System to Facilitate Electrolysis Hydrogen Production. *Int. J. Hydrogen Energy* **2016**, *41*, 8514–8524.
3. Kumar, S.S.; Himabindu, V. Hydrogen Production by PEM Water Electrolysis—A Review. *Mater. Sci. Energy Technol.* **2019**, *2*, 442–454.
4. Smolinka, T.; Bergmann, H.; Garche, J.; et al. Chapter 4—The History of Water Electrolysis from Its Beginnings to the Present. In *Electrochemical Power Sources: Fundamentals, Systems, and Applications*; Elsevier: Amsterdam, Netherlands, 2022; pp. 83–164.
5. Blasi, A.; Fiorenza, G.; Freda, C.; et al. 6-Steam Reforming of Biofuels for the Production of Hydrogen-Rich Gas. In *Membranes for Clean and Renewable Power Applications*; Woodhead Publishing: Cambridge, UK, 2014; pp. 145–181.
6. Schlapbach, L.; Züttel, A. Hydrogen-Storage Materials for Mobile Applications. *Nature* **2001**, *414*, 353–358.
7. Khan, M.A.; Zhao, H.; Zou, W.; et al. Recent Progresses in Electrocatalysts for Water Electrolysis. *Electrochim. Energy Rev.* **2018**, *1*, 483–530.
8. Altinisik, H.; Celebi, C.; Özden, A.; et al. A Review on Membranes for Anion Exchange Membrane Water Electrolyzers. *Renew. Sustain. Energy Rev.* **2026**, *226*, 114530.
9. Khan, A.M.; Nielsen, M.R.; Golsorkhi, M.S.; et al. Multiphysics Modeling of Electrolyzers under Dynamic Converter Operation. *Int. J. Hydrogen Energy* **2025**, *176*, 105–116.
10. Narayanan, A.; Jensen, K.J.; Bisgaard, T.; et al. Multiscale and Multiphysics Simulation Framework for Alkaline Electrolyzer Stacks. In Proceedings of the European Electrolyser & Fuel Cell Forum (EFCF 2025); Lucerne, Switzerland, 1–4 July 2025; p. A0506.
11. Chen, K.; Hu, K.; Yang, J.; et al. Multiphysics Modeling and Parameter Analysis of Industrial Power-to-Hydrogen Electrolyzers. In Proceedings of the 2023 6th International Conference on Energy, Electrical and Power Engineering (CEEPE), Guangzhou, China, 12–14 May 2023.
12. Wang, Y.; Lin, H.-W.; An, W.; et al. Comparison and General Law Research of Multiple Machine-Learning Models for Proton Exchange Membrane Electrolytic Cell Parameters Prediction. *Discov. Appl. Sci.* **2025**, *7*, 546.
13. Xu, B.; Ma, W.; Wu, W.; et al. Degradation Prediction of PEM Water Electrolyzer under Constant and Start-Stop Loads Based on CNN-LSTM. *Energy AI* **2024**, *18*, 100407.
14. Chen, K.; Laghrouche, S.; Djerdir, A. Prognosis of Fuel Cell Degradation under Different Applications Using Wavelet Analysis and Nonlinear Autoregressive Exogenous Neural Network. *Renew. Energy* **2021**, *179*, 802–814.
15. Cheng, Y.; Huang, C.N.; Kuo, J.K.; et al. Optimization of Alkaline Water Electrolyzers through Taguchi Statistical Screening and Grey Wolf Metaheuristics. *Int. J. Hydrogen Energy* **2025**, *192*, 130–142.