

## Smarter Energy Management for Multi-Stack Fuel Cells: AI Coordination for Heavy-Duty Transport

Imagine a freight truck crossing a continent not with one giant hydrogen fuel cell (FC), but with a team of smaller stacks, each quietly playing its part. Some work harder on uphill stretches, others rest when temperatures rise, and a few step in when a teammate begins to weaken. This is the vision behind multi-stack fuel cell hybrid electric vehicles (MFCHEVs): modular systems in which several low-power FC stacks operate together with a battery to deliver traction power. Such architectures promise higher efficiency, improved durability, and enhanced fault tolerance. However, these benefits can only be realized through artificial intelligent (AI) energy management and effective coordination; otherwise, stacks may degrade unevenly, waste hydrogen, or fail prematurely.

This study investigates how AI-based EMS can transform MFCHEVs into coordinated, self-optimizing systems. Managing multiple FC stacks introduces significant complexity, as each operates under varying conditions and degrades at different rates. Conventional rule-based and optimization strategies struggle to cope with these nonlinear and time-varying dynamics, underscoring the necessity of AI-based management. AI enables adaptive, data-driven decision-making capable of learning optimal stack coordination in real time, thereby improving efficiency and extending system lifetime.

The research outlines the technical challenges of managing multi-stack configurations, reviews the evolution of EMS technologies, and identifies common issues such as uneven aging and fault handling. An AI-driven approach based on deep reinforcement learning (DRL) and multi-agent coordination is proposed and compared with traditional strategies. Simulation results demonstrate that information sharing among stacks can enhance performance by up to 61% in degradation-sensitive scenarios. Beyond road vehicles, the proposed framework also shows potential for marine propulsion, rail traction, and stationary hydrogen microgrids, where similar coordination challenges arise.

### 1. Why Multi-Stack Fuel Cells?

Global concerns over fossil fuel depletion and rising pollution have accelerated the transition toward cleaner transportation technologies. In response to increasingly stringent emissions regulations, automakers are investing heavily in electrified powertrains. Among these, fuel cell hybrid electric vehicles (FCHEVs) have emerged as a promising solution, offering high efficiency, fast refueling, and zero tailpipe emissions. By combining hydrogen-powered fuel cells with battery systems, FCHEVs provide both long driving range and strong transient performance, making them particularly suitable for heavy-duty and long-distance applications where full battery electrification remains challenging. Traditionally, heavy-duty FCHEVs have been powered by a single high-power stack. This design is mechanically simpler and easier to control but has significant limitations. A large stack must often operate across a wide range of power demands, from low-load cruising to steep-grade hauling. Outside its optimal load range, efficiency drops and degradation accelerates. Moreover, the entire vehicle's powertrain depends on a single point of failure: if that stack fails, the vehicle is immobilized. Maintenance becomes an all-or-nothing affair, often requiring full replacement of the unit.

Multi-stack FC architectures offer a different philosophy. Instead of one large generator, several smaller FC stacks operate in parallel, supported by a battery. This modularity offers several benefits:

- **Efficiency gains:** Each stack can run closer to its optimal efficiency point, reducing hydrogen consumption.
- **Durability improvements:** Workload can be rotated among stacks to spread wear evenly, prolonging the lifespan of the system.
- **Fault tolerance:** A failed or degraded stack can be isolated while others continue to provide power, enabling limp-home operation.
- **Serviceability:** Modular stacks can be replaced or serviced individually, reducing downtime and cost.

Figure 1 highlights the key challenges associated with designing multi-stack FCs systems. While each of these challenges has important implications, this article focuses specifically on the design of appropriate EMS, due to their critical role in real-time power allocation and control, reducing hydrogen consumption, and prolonging system longevity.



**Figure 1.** General Challenges facing multi-stack FC heavy applications.

## 2. Energy Management Strategy

In any hybrid electric vehicle, the EMS serves as the brain of the powertrain. It decides, at each moment, how much power should come from the FCs and how much from the battery. In multi-stack fuel cell hybrid electric vehicles (MFCHEVs), this role becomes more complicated: instead of one FC source, the EMS must coordinate several, each with its own operating characteristics, health status, and efficiency curve.

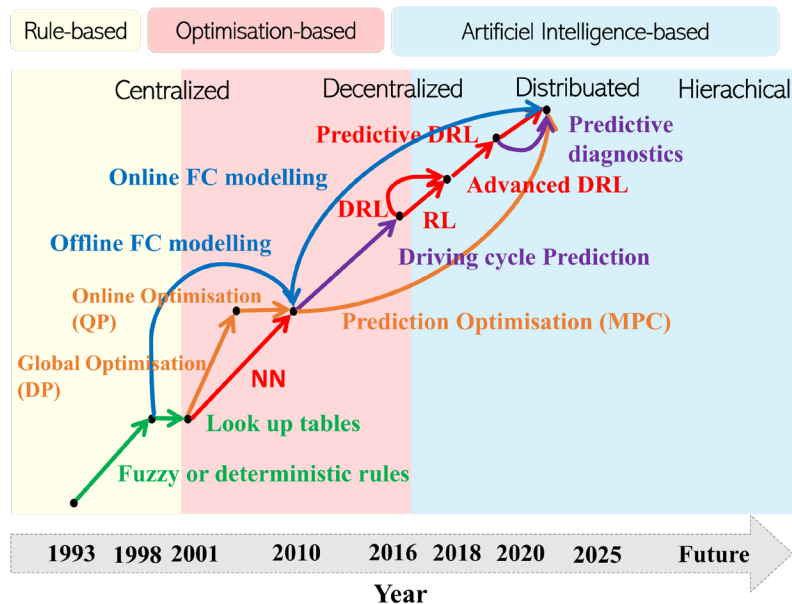
EMSs have evolved in sophistication in parallel with advancements in FC technology. The management of a MFCHEV, in particular, presents a complex control challenge. Such systems are composed of multiple, heterogeneous FC stacks, each with distinct efficiency profiles, performance characteristics, and states of health (SoH). The EMS must dynamically coordinate these individual units, determining the optimal power distribution among them to meet fluctuating power demands in real-time. This coordination requires balancing the dual objectives of maximizing short-term fuel economy while ensuring long-term system durability by minimizing degradation.

Early systems relied heavily on rule-based approaches, fixed if-then-else logic tuned by engineers based on known operating conditions. Rule-based EMSs are straightforward to implement and computationally light, making them attractive for early prototypes. However, they struggle in highly variable or unforeseen scenarios. For instance, if one stack suddenly loses efficiency due to an emerging fault, a fixed rule set may not adapt quickly enough.

The subsequent advancement in EMS development was the introduction of optimization-based strategies, wherein energy management is formulated as a constrained optimization problem. Depending on the application context, the objective function may be defined to minimize fuel consumption, extend stack longevity, or achieve a trade-off between the two over a driving cycle. Techniques such as Global Optimisation Dynamic Programming (DP), which

finds the optimal solution by working backward from a final state, and Pontryagin’s minimum principle (PMP), which minimizes a Hamiltonian function over time, have been applied successfully in single-stack and some dual-stack systems. However, DP requires knowledge of the driving cycle in advance and is not suitable for real-time applications due to its high computational time. These issues have led to the development of more computationally efficient methods, such as Online Optimisation Quadratic Programming (QP) and Model Predictive Control (MPC). An EMS designed with QP minimizes a quadratic cost function (such as the relationship between FC power and efficiency), subject to constraints like battery state of charge and FC power limits. In contrast, MPC is a predictive framework composed of three main components: a predictive model, rolling optimization, and feedback correction. A key advantage of MPC is its flexibility, as the rolling optimization process can employ various optimization techniques to determine the power split based on system complexity. Optimization-based EMSs can find near-optimal solutions under known conditions, but they often require detailed models and may become computationally intensive, especially when scaled to multiple stacks with nonlinear, time-varying behavior.

More recently, attention has shifted to AI-based approaches, which can learn from data and adapt online. Supervised learning can be used to predict efficiency or degradation trends from operational data, while unsupervised methods can identify patterns in degradation modes without explicit labels. Reinforcement learning (RL) and its deep learning variant (DRL) are particularly appealing because they can learn optimal control policies through interaction, without needing an exact model of the system. This is valuable in multi-stack systems, where modeling every detail is impractical. While AI offers flexibility and adaptability, it also presents challenges in training data requirements, safety assurance, and interpretability. Nevertheless, for MFCHEVs where complexity, uncertainty, and adaptation demands are high, AI is increasingly seen as the most promising path forward. Figure 2, drawn from the literature, illustrates the chronological evolution of algorithms applied in FC systems. In recent years, the advent of RL has introduced a new level of adaptability and responsiveness to FC control. In order to identify the gaps in designing EMS for multi-stack FC systems, next section briefly reviews the literature.



**Figure 2.** Energy management development history in FCEVs and key milestones from 1993 to 2025.

### 3. Synopsis of the Literature

The growing interest in sustainable transportation has intensified research into advanced powertrains, particularly those utilizing multiple FC stacks in a hybrid configuration. The foundational appeal of this architecture lies in the

principle of modular energy systems, which leverages independent, standardized power sources to enhance economic and technical performance. As Soltani et al. explain, this modularity helps overcome the inherent limitations of cost, durability, and efficiency associated with single-source systems, providing a strong justification for multi-stack designs. However, the primary challenge of this approach is the development of a sophisticated EMS capable of optimally coordinating these multiple power sources. Early and foundational approaches to this challenge often relied on established control strategies. For instance, Jiang et al. applied an Equivalent Consumption Minimization Strategy (ECMS) to a multi-stack hydrogen electric train, focusing primarily on improving the long-term durability of the fuel cell system by managing the power split with a lithium-ion battery. These methods establish a strong baseline for performance but can lack the adaptability of more dynamic, learning-based systems.

To address this, recent research has increasingly shifted toward machine learning, particularly RL. Ghaderi et al. introduced an RL-based EMS that learns to manage power distribution in real-time, effectively minimizing the combined operational costs of hydrogen consumption and component degradation. Building on this, Shi W et al. employed a Hybrid Quantum Reinforcement Learning algorithm, which achieved superior economic performance with a less complex model. Further exemplifying this trend, Xiaohong Yuan et al. developed a two-layer EMS based on the Deep Deterministic Policy Gradient (DDPG) algorithm for a large, six-stack hybrid tractor. Their work, validated through both simulations and on-road experiments, also introduced a novel three-group stack architecture that demonstrated superior stability and lower power loss. These learning-based approaches allow the EMS to adapt to varying driving conditions without pre-defined rules. The latest advancements focus on tackling the challenges of scalability and real-time processing in increasingly complex systems. Khalatbarisoltani A et al. presented a cutting-edge decentralized EMS that combines Decentralized Model Predictive Control with Federated Reinforcement Learning. This hybrid methodology creates a fully scalable system that outperforms traditional centralized methods in precision, convergence speed, and real-time applicability. This trend towards intelligent, decentralized control, alongside novel powertrain configurations like the dual-fuel-cell design proposed by Eom B et al. for enhanced efficiency and robustness, marks the frontier of research in optimizing MFCHEVs.

## **4. Key Challenges in Multi-Stack FCs EMS**

### **4.1 When One Stack Ages Faster**

One of the most insidious problems in multi-stack systems based on the literature is uneven degradation. Even if stacks are nominally identical, small differences in manufacturing, cooling, or gas distribution can cause them to age at different rates. Over time, an imbalance emerges: some stacks retain near-new performance while others show reduced voltage output or increased internal resistance. If the EMS does not account for this, it may inadvertently overuse the healthier stacks since they appear more efficient while sparing the weaker ones. This accelerates the imbalance, leading to early retirement of the stronger stacks and underutilization of the weaker ones. The end result is reduced overall system life.

Addressing this requires health-aware power allocation, where the EMS explicitly considers SoH in its decision-making. The challenge is that SoH is not directly measurable; it must be inferred from voltage-current behavior, impedance spectroscopy, or long-term performance trends. AI can help by learning degradation signatures from historical data and predicting future health trajectories.

### **4.2 Operating After a Stack Fault**

FCs can fail in various ways: sudden membrane rupture, gas leakage, catalyst poisoning, or cooling system breakdown. In a single-stack system, any such fault typically ends the journey. In a multi-stack system, however, the EMS can isolate the faulty unit and redistribute load among the remaining stacks. Doing this well requires fault-tolerant EMS design. The system must detect the fault quickly, assess the remaining power capacity, and adjust the operating strategy to meet demand without overloading the healthy stacks. AI-based EMSs can be trained on simulated fault scenarios, enabling them to respond appropriately to rare but high-impact events.

### 4.3 Coordination between stacks

Poor coordination can cause imbalanced power distribution, leading to higher degradation and lower efficiency. Coordination strategies in MFCHEVs typically fall into three categories: cooperation, independence, and information sharing, each with trade-offs in scalability, reliability, and complexity. While cooperation can improve efficiency and maintain operation when some stacks underperform, it may also cause new inefficiencies if stacks lack awareness of each other's health and operating conditions.

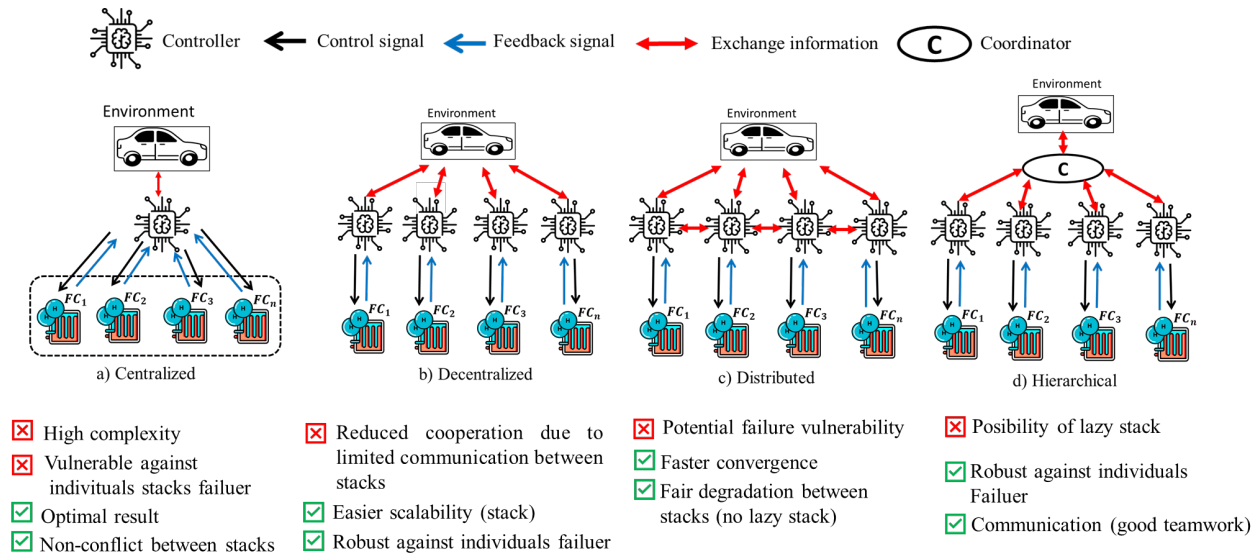
Information sharing between stacks allows for better power balancing, reduces uneven degradation, and improves reliability by enabling coordinated operation and more uniform wear over time. This prevents certain stacks from being overused while others remain idle. However, complete reliance on cooperation or communication can reduce system robustness. Maintaining a degree of independence is crucial to ensure fault tolerance, preventing failures in one stack from affecting others and enabling continued operation in case of partial system failure.

### 5. Coordinating Multiple Stacks

A clear understanding of coordination mechanisms in multi-stack FC systems is essential to address the identified gaps in EMS design. In this context, each stack can be viewed as an autonomous agent that communicates and collaborates with others to achieve global objectives such as energy efficiency and load balancing. Four main coordination architectures are commonly observed:

- **Centralized:** One controller makes all decisions. Simple to design but vulnerable to a single point of failure.
- **Decentralized:** Each stack operates independently based on local measurements. Robust to controller faults but risks suboptimal global performance.
- **Distributed:** Local controllers make decisions while sharing information with neighbors or a central node. Balances robustness and coordination.
- **Hierarchical:** A high-level coordinator sets goals or constraints for local controllers, combining centralized oversight with decentralized execution.

For MFCHEVs, distributed and hierarchical architectures are often the best compromise, allowing local autonomy while enabling global optimization. Figure 3 shows these architectures.



**Figure 3.** Classifications of the coordination methods in MFCHEVs.

## 6. AI as an Energy Management

The coordination of a multi-stack FC system requires a sophisticated EMS. While each stack can operate as an independent power unit, the EMS governs their collective behavior to ensure optimal load distribution, dynamic response, and overall system stability. However, as the number of stacks increases, so does the system's complexity interactions among stacks become nonlinear, performance characteristics drift with aging, and real-time environmental variations make it difficult to maintain efficiency through static control rules. Traditional rule-based or optimization-based strategies often lack the adaptability and scalability required to manage these interdependent, time-varying processes effectively.

In this context, artificial intelligence (AI), particularly through multi-agent systems, offers a natural and powerful framework for control and coordination. AI enables multiple intelligent agents each representing an FC stack or subsystem to collaborate, share information, and make autonomous yet coordinated decisions that balance individual health and global system objectives. This cooperative intelligence allows for real-time adaptability, fault tolerance, and self-optimization, which are crucial for ensuring long-term durability and efficiency in complex hybrid architectures. The ability to dynamically allocate power and respond to degradation patterns makes AI indispensable for managing multi-stack FC systems, especially under unpredictable driving conditions.

There are various types of AI algorithms, such as supervised learning, unsupervised learning, and RL, increasingly applied in energy management where real-time adaptability is essential. Supervised learning has been widely adopted for tasks such as performance prediction, fault detection, and health estimation. While supervised learning is effective when labeled datasets are available and the goal is to predict specific outcomes, unsupervised learning is valuable when labeled data is scarce or unavailable, as it enables the exploration of complex, high-dimensional datasets to extract meaningful insights without prior knowledge of system outputs. RL, often regarded as a semi-supervised approach, allows an agent to learn optimal behavior by maximizing cumulative rewards through interactions with the environment, even in the absence of explicitly labeled data. In recent years, RL has evolved into DRL by integrating deep neural networks, enabling efficient feature extraction from high-dimensional state and action spaces. This allows it to solve problems that traditional RL methods struggle to handle.

## 7. Proposed Hybrid Coordination

To address some of the identified gaps in EMS design for MFCHEVs, this article proposes a DRL-based coordination strategy aimed at enhancing adaptability, coordination efficiency, and overall system performance. Leveraging the capabilities of AI in multi-agent systems, a hybrid hierarchical–distributed AI-based EMS is introduced, where each FC stack is managed by its own intelligent agent. This architecture enables autonomous local decision-making while maintaining global coordination through the sharing of each stack's policy. The primary objective is to minimize the system's total operational cost by fostering cooperative behavior among stacks.

To mitigate the risk of “lazy” stacks, defined as units contributing less to the power demand, an information-sharing mechanism is implemented, in which stacks periodically communicate with a centralized aggregator. The aggregator evaluates individual stack performance and coordinates power allocation accordingly.

To improve robustness, enabling real-time adaptability requires integrating online FC modeling, which continuously updates each stack's parameters based on its operating conditions, thereby capturing dynamic changes such as performance decline or efficiency loss.

The reward function is designed to balance hydrogen efficiency, battery usage, and stack health preservation, while imposing penalties for operating outside safe limits.

## 8. DRL in Action: Simulation Results

To evaluate the proposed DRL-based multi-stack coordination strategy, we conducted a series of high-fidelity simulations using a validated heavy-duty truck model. The model included three 500 W (scaled up to 7 kW) low-temperature proton exchange membrane (PEM) FC stacks, a 4 kW lithium-ion battery, and an electric drive system rated for 28 kW peak output. Ambient temperature variations, realistic road gradients, and transient traffic conditions were integrated to mimic operational variability.

To meet the power requirements of the three-wheel vehicle (28 kW and 96 V) while utilizing available experimental data, a scaling procedure was implemented. The output power of a 500-W FC system was linearly scaled up to 7.33 kW. This linear scaling is a valid functional simplification because the core optimization properties of the energy management problem are unaffected. The characteristic shapes of the polarization, power, and efficiency curves are preserved during scaling, ensuring that an optimization algorithm capable of finding the optimal operating points on a 500-W system can do so on a scaled system as well. It is important to acknowledge the limitations of this approach. A physical 7.33 kW system would require a significantly more complex Balance of Plant (BoP) (e.g., a liquid-cooling system) than the 500W air-cooled experimental unit. Furthermore, this more complex BoP would likely result in a slower dynamic response, a factor that is relevant to stack degradation. However, since the primary contribution of this work is the design and comparison of high-level EMS, this simplification allows for a valid assessment of the algorithms' ability to manage power distribution and degradation based on the core efficiency characteristics.

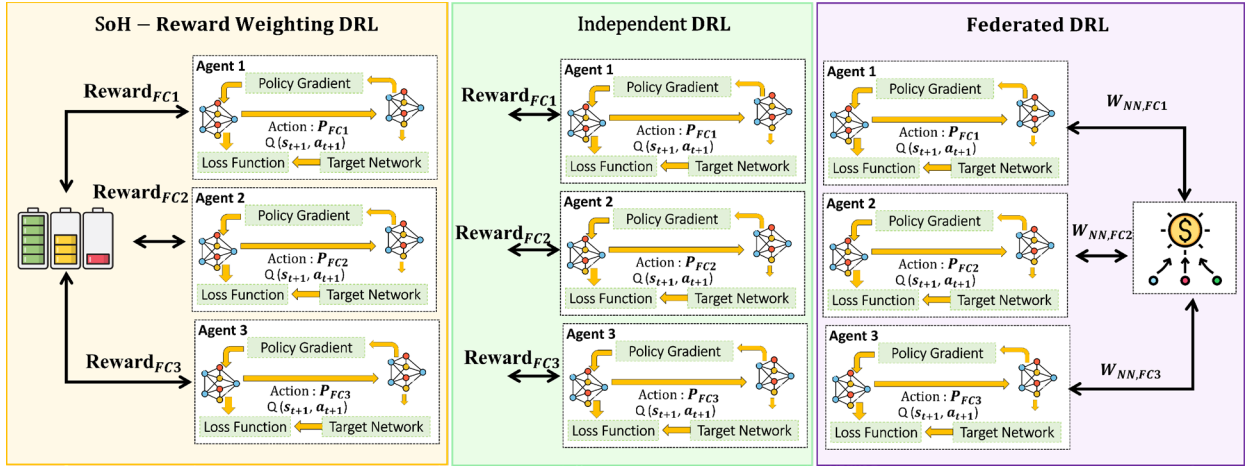
The considered reward function is formulated as the sum of the Hydrogen Cost, FC Degradation Cost, and Battery Degradation Cost. Details on the sizing and characteristics of the power sources and vehicle configuration, and the formulation of hydrogen consumption, degradation, and the reward function are provided in the For Further Reading section.

### **8.1 Simulation Setup**

The EMSs evaluated in this study include the following architectures:

1. Independent DRL (IDRL): Each stack trained its own agent using only local measurements (stack voltage, current, temperature, estimated SoH). No inter-stack communication was allowed.
2. Federated DRL (FDRL): Following an approach similar to IDRL, each stack began by training its own agent on local data. However, to foster collaboration, a central aggregator then collected and averaged the agents' policies. This updated, averaged policy was redistributed to all stacks, which prevents policy divergence and ensures consistent, collective improvement.
3. SoH-Reward Weighting DRL (SWDRL): Similar to FDRL but with weighting in the reward function to discourage overuse of stacks with lower SoH.

These architectures are depicted in Figure 4.



**Figure 4.** Architecture of the Implemented DRL Methods.

The training driving cycle should be specifically designed to ensure the resulting policies are robust and generalizable across a wide range of real-world conditions. For this purpose, this paper utilized the Worldwide Harmonized Light Vehicles Test Cycle (WLTC) Class 3b. The WLTC is a globally recognized standard, designed to reflect realistic driving conditions, and comprises the following conditions:

- Urban Delivery: stop–start cycles with frequent low-speed operation and rapid accelerations.
- Highway Freight: sustained high load with moderate transients.
- Mixed Long-Haul: realistic profile combining long climbs, descents, and urban segments.

Each simulation ran for a virtual distance of 1800 s, with degradation models accounting for electrochemical aging and mechanical wear.

All three strategies rely on agents that observe the battery's SoC and the vehicle's power demand ( $P_{req}$ ) to control the power output ( $P_{FC}$ ) of their respective FC stacks. The key difference lies in their coordination. The baseline IDRL uses no coordination, each agent acts as an isolated decision-maker. Building upon this, FDRL introduces a communication layer where agents periodically share their learned neural network parameters, which are averaged into a global model and redistributed to enhance collective intelligence. The SWDRL refines this process further by using the SoH of each stack as a weight during the averaging, ensuring that healthier, more reliable stacks have a greater influence on the shared strategy.

## 8.2 Key Findings

The simulation outcomes provide critical insights into the performance, degradation management, and coordination of the multi-stack FC system under different EMSs. To comprehensively evaluate the effectiveness of the proposed EMS, we assessed the coordination between FCs by analyzing their SoH-based Degradation Awareness. This metric reflects the DRL agent's ability to manage each FC's power output effectively, with the goal of minimizing factors that accelerate degradation. It is quantitatively measured using the Average Standard Deviation of FC Power Output, where a lower standard deviation indicates more stable and consistent power distribution, reduced stress on the FCs, and a slower rate of degradation. Thus, a lower Average FC Power Standard Deviation corresponds to better Degradation Awareness.

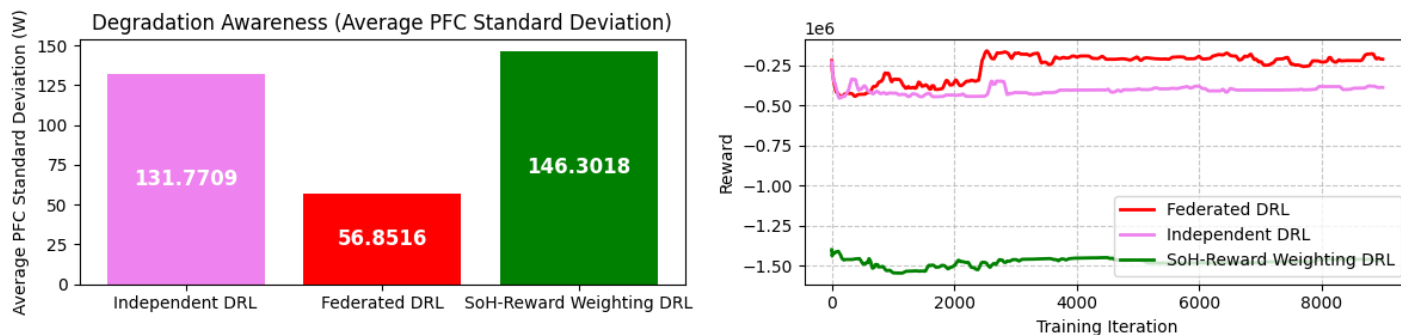
Figure 5 illustrates both the reward convergence and degradation awareness for the implemented DRL methods. As observed, the FDRL approach achieves a higher reward convergence. However, when the reward is weighted by the

SoH of the FCs, its value decreases, indicating higher system cost (Hydrogen Cost, FC Degradation Cost, and Battery Degradation Cost) and suboptimal DRL performance. The figure also shows that the degradation awareness in FDRL reaches approximately 56%, which is 61% higher than that achieved by IDRL and SWDRL.

The results highlight the value of information sharing in multi-stack systems. While IDRL is adaptive, it lacks the situational awareness required for effective coordination. FDRL demonstrates that global optimization enhances both reward and degradation awareness. In contrast, SWDRL, which incorporates explicit health weighting in the reward, proved less effective than FDRL and caused the reward to converge at a lower level.

The different performance of SWDRL stems directly from the conflict created by its reward structure. Unlike FDRL, which learns to manage system health implicitly as part of a holistic, long-term strategy, SWDRL imposes a rigid, explicit penalty for using stacks with lower SoH. This forces the FCs to learn a suboptimal and overly cautious policy, often making inefficient decisions like running healthy stacks in poor operating ranges just to avoid the immediate health penalty. This consistent, inefficient behavior prevents the system from reaching its true performance potential, causing the agent's cumulative reward to converge at a lower level. This weighting approach acts as a handbrake, preventing the system from discovering the optimal trade-off between short-term efficiency and long-term degradation a balance the more flexible FDRL model achieves through global information sharing.

The differences also reveal that short-term optimality does not guarantee long-term sustainability. Strategies that focus solely on hydrogen efficiency may inadvertently concentrate stress on particular stacks, accelerating their degradation. By contrast, health-aware strategies ensure that all stacks reach end-of-life at similar times, maximizing the return on investment for the entire powertrain.



**Figure 5.** Comparison of reward and degradation awareness of implemented DRL methods.

## 9. Conclusion and Outlook

Multi-stack fuel cell hybrid electric vehicles (MFCHEVs) represent a significant step toward robust, efficient, and maintainable zero-emission heavy-duty transport. Their modular architecture offers inherent redundancy, operational flexibility, and the potential for finely tuned efficiency optimization. However, realizing these benefits depends critically on the energy management strategy's ability to balance conflicting objectives in real time.

Our simulations demonstrate that AI-based coordination, particularly DRL with shared policy information, can deliver tangible gains in both reward and degradation awareness, thereby preventing lazy stacks and improving overall system performance. These findings suggest that in mission-critical applications, where system longevity and availability are paramount, such health-aware strategies may be the preferred choice.

The implications of these findings extend well beyond road freight, offering insights that are relevant to the development and optimization of advanced energy management technologies across a variety of transportation and power systems.

- Marine propulsion: Ferries and coastal vessels with modular fuel cell (FC) racks can dynamically allocate power based on stack health, reducing maintenance downtime in port.
- Rail traction: Modular FC pods on hydrogen trains can remain in service despite individual module degradation, improving fleet availability.
- Stationary microgrids: Off-grid hydrogen power systems can optimize stack rotation to meet seasonal load variations while preserving component life.

Looking ahead, the next generation of multi-stack FCs EMS will likely integrate:

1. Predictive degradation modeling to anticipate failures before they affect performance.
2. Cyber-secure distributed coordination, ensuring resilience against communication faults or malicious interference.
3. Route- and weather-aware optimization, where EMS decisions adapt to predicted terrain, traffic, and ambient conditions.
4. Human-AI collaboration, where operators can guide or override AI decisions with minimal training, ensuring trust and safety in high-value fleets.

Ultimately, the path forward for MFCHEVs is not just about hardware modularity, it is about turning a set of independent FC stacks into a cooperative, self-optimizing energy ecosystem. Achieving that requires AI systems that think not only in seconds, but also in years, aligning day-to-day efficiency with the long-term health of the powertrain. If done well, the payoff will be fleets that run cleaner, longer, and smarter.

## 10. For Further Reading

Ghaderi, R., Kandidayeni, M., Boulon, L. and Trovao, J.P., 2023. Q-learning based energy management strategy for a hybrid multi-stack fuel cell system considering degradation. *Energy Conversion and Management*, 293, p.117524.

Shi, W., Sun, X., Zhang, Z., Chen, J., Du, Y., Ruan, J., Ding, Y., Wang, L., Huangfu, Y. and Xu, Z., 2025. Optimal energy management for multi-stack fuel cell vehicles based on hybrid quantum reinforcement learning. *IEEE Transactions on Transportation Electrification*.

Adnane, M., Nguyễn, B.H., Khoumsi, A. and Trovão, J.P.F., 2025. Real-time Multi-Ensemble Learning-Based Power Prediction for Energy Management of Dual-Source Electric Vehicles. *IEEE Transactions on Vehicular Technology*.

Soltani, A.K., Kandidayeni, M., Boulon, L. and St-Pierre, D.L., 2019. Modular energy systems in vehicular applications. *Energy Procedia*, 162, pp.14-23.

Khalatbarisoltani, A., Boulon, L. and Hu, X., 2023. Integrating model predictive control with federated reinforcement learning for decentralized energy management of fuel cell vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 24(12), pp.13639-13653.

Shi, W., Sun, X., Zhang, Z., Chen, J., Du, Y., Ruan, J., Ding, Y., Wang, L., Huangfu, Y. and Xu, Z., 2025. Optimal energy management for multi-stack fuel cell vehicles based on hybrid quantum reinforcement learning. *IEEE Transactions on Transportation Electrification*.

Eom, B., Eom, K., Yang, D. and Kim, M., 2024. Novel electric vehicle powertrain of multi-stack fuel cell using optimal energy management strategy. *International Journal of Automotive Technology*, 25(2), pp.201-211.

Jiang, D., Long, Y., Fu, P., Guo, C., Tang, Y. and Huang, H., 2024. A novel multi-stack fuel cell hybrid system energy management strategy for improving the fuel cell durability of the hydrogen electric Multiple Units. *International Journal of Green Energy*, 21(8), pp.1766-1775.

Yuan, X., Ma, Z., Wen, S., Liu, X., Su, C. and Wang, Y., 2025. Energy management strategy for a hybrid six-stack fuel cell vehicle based on DDPG. *Energy*, p.138606.

## Biographies

**Razieh Ghaderi** (razi.ghaderi@usherbrooke.ca) is with the Department of Electrical and Computer Engineering, University of Sherbrooke, Sherbrooke, QC J1K 2R1, Canada, and Université du Québec à Trois-Rivières, Trois-Rivières, QC G8Z 4M3, Canada.

**Loïc Boulon** (loic.boulon@uqtr.ca) is with the Department of Electrical and Computer Engineering, Université du Québec à Trois-Rivières, Trois-Rivières, QC G8Z 4M3, Canada.

**João Pedro F. Trovão** (joao.trovao@usherbrooke.ca) is with the Department of Electrical and Computer Engineering, University of Sherbrooke, Sherbrooke, QC J1K 2R1, Canada, and IPC-ISEC and INESC Coimbra, 3030-199 Coimbra, Portugal.

**Minh C Ta** (cao.minh.ta@usherbrooke.ca) is with the Department of Electrical and Computer Engineering, University of Sherbrooke, Sherbrooke, QC J1K 2R1, Canada.