Engineering and Technology Journal e-ISSN: 2456-3358

Volume 11 Issue 02 February-2026, Page No.-9034-9049

DOI: 10.47191/etj/v11i02.29, I.F. – 8.482

Data-Driven Predictive Diagnostics and Fault-Tolerant Control for Fuel-Cell Electric Vehicle Powertrains Under Uncertainty

Adel Elgammal

Professor, Utilities and Sustainable Engineering, The University of Trinidad & Tobago UTT

ABSTRACT: Fuel-cell electric vehicles (FCEVs) are susceptible to sensor drift, actuator degradation and air path/water management faults that can quickly deteriorate efficiency, drivability and durability. This work introduces a control-oriented data-driven framework to integrate predictive diagnostics with FTC for FCEV powertrains subject to uncertainty. A multi-sensor feature pipeline comprising stack voltage/current, cathode pressure, compressor speed, hydrogen flow rate DC-link voltage and traction power is first established and the health indicators are generated via a probabilistic sequence model modeling temporal dependencies. The diagnostic module conducts an online detection and isolation of the faults in early stages with fault magnitude estimation for remaining useful time to predict constraint violation. Second, an active FTC layer allows real-time fault estimation–based control reconfiguration to update an uncertainty-aware powertrain controller by which fuel-cell ramp-rate limits, air-path constraints and energy-storage limits are enforced. The controller ensures that fast power transients are absorbed by the battery/supercapacitor, while keeping the fuel-cell system operational at high efficiencies and ensuring traction power tracking in addition to safe limits during faults. Uncertainty is dealt with by disturbance modeling and confidence-weighted adaptation to avoid overreacting to noisy diagnostic outputs. Simulation results for representative urban and aggressive drive cycles are presented, which incorporate introduced sensor bias, compressor efficiency loss and hydrogen starvation scenarios that show superior robustness compared with the non-tolerant 'baseline' in terms of increased fault detection lead time, no constraint or emissions coverage issues as well as efficient minimisation of hydrogen consumption while maintaining drivability. The considered combination of predictive diagnostics and FTC can thus offer a practical approach to reliable FCEV operation that will bring quantifiable safety and efficiency gains under real-world automotive scenarios.

KEYWORDS: Fault-tolerant control; Predictive diagnostics; Fuel-cell electric vehicles; Data-driven fault detection and isolation (FDI); Uncertainty-aware control reconfiguration.

INTRODUCTION

Proton exchange membrane fuel cell (PEMFC) based fuel-cell electric vehicles (FCEVs) have been considered as a promising option for low-carbon transportation due to their characteristics of fast refueling, long driving range, and zero tailpipe emissions. In real automotive powertrains, however, the PEMFC stack is only a component of closely interrelated systems that also involve the balance-of-plant (BoP) subsystems (air compressor, humidifier, valves, pumps, thermal loop), hydrogen store and recirculation system HV distribution and power electronics (DC/DC converters and traction inverter) [5]. Most FCEVs are also equipped with an energy storage stage (battery and/or supercapacitor) to meet fast traction requests while protecting the fuel cell against aggressive transients, for example during regenerative braking [23], [24]. This multi-source and multi-domain coupling gives complex nonlinear behavior, several time scales and hard operational constraints that have to be fulfilled to ensure the safety, efficiency and durability [5], [6]. As a consequence, FCEV powertrains control is ideally hierarchical. A supervisory energy management system coordinates demanded power between the fuel cell and storage systems whereas a set of lower-level controllers control air-path variables (e.g., excess oxygen ratio), compressor speed, manifold pressure, and DC-bus voltage [5]. Constraints are primarily focused at: the fuel cell (ramp rates, and oxygen/air path restrictions), the battery and supercapacitor (voltage/SOC/current bounds), and the inverter and traction motor (torque–power boundary) [10], [11]. Similar limitations apply to regenerative torques and charge acceptance under braking, where a mixture of regenerative and frictional braking is necessary in order to achieve the desired deceleration while preserving stability and driveability [23]. One can observe that the interrelation between traction strategy and subsequently possible regenerative capability (via SOC, or voltage headroom6) is particularly significant in a transitory driving cycle because of the number of stop–go events, which may lead to full store saturation if storage peaks are directly interfaced by regeneration [24]. One of the most persistent hurdles to large-scale implementation is reliability. PEMFC systems show a wide range of fault types whose list is: (i) stack/internal such as flooding, drying, starvation and leakages;

(ii) BoP such as compressor degradation or surge, valve sticking, humidification issues and cooling loop anomalies; (iii) sensor related like bias/drift in current/pressure/temperature/flow sensors; (iv) HV / power-electronics related for insulation degradation and converter faults [5], [6]. These faults may be slow or sudden, single or simultaneous; and commonly appear indirectly through coupled dynamics rather than a distinct, isolated signature [4], [16]. The health-management review for PEMFCs of electric vehicles stressed that fault modes, diagnosis and mitigation should be treated together since the operation decision affects degradation paths and vice-versa [5]. Classical fault diagnosis works on PEMFC systems are widely performed by model-based approaches including analytical redundancy and residual generation, where the derived consistency relations from physics-based or identified models are compared with sensor measurements for the purpose of fault detection and isolation (FDI) [2]. Model-based PEMFC diagnosis is the focus of pioneering research work explaining how residuals can be used in FDI and emphasizing that characterization of residual sensitivity at different operating points and under disturbances is a crucial requirement [2]. Recent critical reviews and analyses of analytical redundancy based PEMFC diagnosis confirm that modelling error robustness remains a challenge and practical uptake depends on tuning the trade-off between model fidelity, implementability and false alarm rejection [16]. Meanwhile, model-free and data-driven methods (such as neural networks, fuzzy logic, statistical learning and signal processing) have also been extensively discussed as good alternatives when precise diagnosis-oriented models are difficult to establish or when faults are strongly nonlinear and operating-point dependent [3], [4]. More recently, the increase in onboard sensing and fleet data has allowed for an accelerated development of data-driven diagnostics and prognostics. Deep learning was employed for PEMFC water-management fault diagnosis through the combination of segmented-cell measurements and inference, which can provide richer observability in the system internal state than bulk-stack measurements alone [7]. Fleet-level analysis has investigated the trajectory of insulation resistance, and proposed advanced time-series modelling (e.g., LSTM-based predictors) for anomaly detection associated with HV subsystem... and suggested... the importance of continuous data in AD deployment [8]. Ensemble learning to enhance diagnostic robustness and multi-stage workflow that integrates physics-relevant features with machine learning classifiers are among potential complementary methodologies [17], [18]. These trends drive architectures that integrate predictive diagnosis (what is developing, how bad and when it matters) with control-oriented FTC (how to redistribute authority, reshape constraints and keep the vehicle safe and efficient). A fault-tolerant control is the one that transfigures "fault aware" into "safe and performing". Full FTC reviews for PEMFC systems including passive and active FTC were reported, it is demonstrated that any successful method should preserve stability, satisfy constraints and handle both uncertainty as well as faults uncertainties in [1]. Active FTC approaches have been observer-based fault recovery and control reconfiguration, as well as fault accommodating ideas on virtual actuators that enables a nominal controller to work as though the plant is non-faulty [12], [13], [15]. In PEMFC application, Takagi–Sugeno and quasi-LPV virtual actuator systems have been studied to deal with nonlinearities and to maintain the closed-loop behavior in presence of faults [12], [13]. Constrained MPC is also appealing since it naturally handles multivariable constraints; reference-governor MPC for automotive fuel cells and low-complexity formulations of MPC suggest feasibility of real-time implementation [10], [11]. At a system-level scale, the laws developed for PHM [5] indicate that diagnostics and prognostics must be linked to an "action phase" in which health knowledge influences operational decisions but this integration is mature to a lesser extent in most of the PEMFC works [6].

Even though attempts have been made in data-driven diagnostics and constraint-based control, most FCEV approaches tend to consider the diagnosis/prognosis and the control as a loosely couple module. Usually, alarms, derating and limp-home logics are activated upon diagnosis of faults, at which point the nominal energy management and air path regulation hardly differs from that under healthy operating conditions [1], [5]. This decoupling is not complying with the nature of many safety- and durability-critical constraints in PEMFC powertrain, which are fault-dependent and time-varying ones even continuous [6]. For instance, the compressor wear or surge likelihood alters controlable dynamic O2 supply; valve sticking affects manifold pressure controllability performance; sensor bias induces systematic constraint violations; and insulation degradation tightens allowable operating HV range [7], [8]. In such a situation, the proper response is not "reducing power," but rather re- formulating constraints and reallocating control authority according to the fault type, severity and confidence [1], [10]. The second dimension of this paradox is a more temporal one. A broad range of PEMFC error processes change slowly over time (e.g., compressor efficiency gradual loss, generation of increasing leaks, membrane hydration imbalance or insulation resistance degrade slow- down), creating a time window for preventive control actions to be taken by sensing and predicting such faults [6], [8]. However, early-stage signatures can be subtle, depend on operating point and are masked by driving transients; and data-driven predictions come with uncertainty and risk suffering from domain shift [17]. In the context of a safety-critical powertrain, such uncertain diagnostic output must be mapped to control action that will remain robust — otherwise the controller might "chase noise", and deterioration drivability, or result in new violations [48]. Therefore, the core problem becomes: How can we develop an integrated architecture that marries data driven predictive diagnostics (FDI + severity + prediction with uncertainty) with control-intensive, constraint-aware fault-tolerant control to ensure safety, drivability and efficiency under uncertainty? [1], [6], [10], [21] PEMFC powerplant pair fast electrical (or convert) dynamics with slower air/hydrogen delivery, humidification and thermal-water processes [5]. Numerous internal parameters, determining safety and degradation, such as local oxygen concentration, catalyst-layer water saturation and the distribution of membrane hydration are not directly measured in production vehicles [3], [5]. Model-based diagnosis produce residuals by utilising measured signals, however its robustness is challenged as new operating condition emerge quickly and model parameters are aging-affected [2], [16]. The classical philosophy of model-based diagnosis on PEMFC systems is demonstrable though not tested to compare residual sensitivity across the regimes, and in the presence of disturbances and modeling errors [2]. Modern reviews also stress the point that AR methods can be brittle in the face of uncertainty, unless used together with robust design and adequate post-processing (thresholding) [16]. Actual FCEV faults are nonuniform and interrelated: stack water management faults can be confused with air-path faltering; compressor wear modifies oxygen supply and potentially causes stress to the stack acceleration; sensor offsets may mislead the controller into running selector sections of the stack in unsafe areas; High Voltage (HV) insulation damage limits acceptable voltage operation [5], [8]. Most methods are designed for one fault but multiple faults often occur and affect the accuracy of isolation [25]. Ensemble learning and multi-stage diagnosis techniques seek to enhance robustness and multi-fault treatment, but at the cost of increasing complexity in feature construction and varying [17], [18]. Both model- and non-model-based PEMFC diagnosis surveys and more general PEMFC diagnosis reviews indicate that signatures of faults are in many cases combined, such as an unambiguous isolation is not possible without a richer sensing or hybrid model/data fusion [3], [4]. For diagnostics based on data, they depend on representative datasets covering drive cycles, ambient conditions, component aging and severity of fault. In reality, labeled fault data are not common: because faults are rare and it is hard to determine when they occur; the root-cause also has various uncertainties [6], [8]. Fleet data sets are also often private, thereby restricting public benchmarking and reproduction [8]. Furthermore, transients have been observed to be similar with fault signatures; enabling techniques that aim at explicitly reducing false alarms in switching regimes [21]. These facts make it difficult to ensure generalisation across platforms, climates and sensor suites. In an FCEV vehicle, fault tolerance is the ability to ∣ Maintain constraint satisfaction and acceptable tracking performance under faults [1], rather than simply detecting anomaly. Under faults and uncertainty, the fuel-cell safety margins (air stoichiometry, pressure, humidification, thermal bounds), ramp-rate limits, and energy storage constraints can all decrease [1], [10]. Optimal MPC approaches, including RGC, have been considered as they can handle multivariable constraints and accommodate reconfiguration through updating the online constraint sets, and automotive implementations indicate real-time capability [10], [11]. However, integrating the MPC with uncertain diagnostic outputs is no trivial matter: fault estimation may be delayed, probabilistic and biased or computational budgets on automotive ECUs remain tight [6], [10]. In fact, PHM workflows often result in health indicators and predictions, but to a lesser extent the loop is closed with operational decisions prescribed that explicitly account for constraints, uncertainty and trade-offs of multiple objects [6]. Work aimed at the action of phase reviews, also points out that many works are still tilted to the observation and analysis rather than decision-making and control integration [6]. Digital-twin and RUL-forecast methods can predict health trajectories across different conditions rather promisingly; however, realizing those predictions in real-time fault-tolerant decisions with guarantee on feasibility is still at early development [9].

To fill these gaps, the key concept of the proposed paper is to couple predictive diagnostics with constraint-aware fault-tolerant control in a tight manner by employing uncertainty as common interface between them. The desired contribution is not the mere improvement of diagnostic accuracy and/or control performance -- although it will certainly be important -- but that a diagnosis prediction-driven approach generates altered control constraints and objectives facilitating feasible solutions that enhance safety, robustness and efficiency concurrently. Diagnostics is posed as a multi-sensor time-series inference task to neurons) that maps stack voltage/current, cathode pressure, compressor speed, hydrogen flow and DC-link voltage (as well as traction power) to the (i) probabilities of different class of faults (FDI), (ii) continuous estimates representing severity – defined as fault magnitude – and in particulars how likely some predictive metrics are violated or fall outside validated operational practitioner-specified levels in the near future with calibrated uncertainty [7], [8], [17]. Studies like deep learning based PEMFC diagnosis inspire the potential of temporal models for learning subtle degradation and water management signatures, while fleet-level insulation anomalies motivate predictive monitoring derived from longitudinal vehicle data [7], [8]. Motivated by ensemble learning and multistage frameworks, in [17], [18], [21], robustness-promoting techniques can be applied for generalization improvement and false alarm reduction. When a fault is detected or diagnosed, the controller does not shift to fixed derating logic, but rather makes reconfiguration of decisions according to diagnosed failure type and seriousness. For example:

- Air-path/BoP faults (such as compressor deterioration/surge risk) are associated with a straightening of fuel-cell ramp constraints and oxygen safety margins, instead while diverting fast transients to storage to avoid starvation or surge pullouts [1], [10].
- Sensor errors justify fault-tolerance through observers / virtual sensors to hide or compensate improper signals pushing biased values towards never reaching an unsafe command as the use of model-based diagnosis focuses us on analytical redundancy and robust residual design [2], [16].
- Site level HV insulation anomalous operation invokes safe envelope management (using storage buffering) and power

shifting, as motivated by the fleet-wide insulation monitoring [8].

MPC (or reference-governor MPC) is a clear candidate, since it can handle multivariable constraints and embed reconfiguration in the form of online constraint adaptation [10], [11]. Complementary fault-hiding controllers based on TS/qLPV actuator virtualization enable the preservation of a nominal controller while addressing faulty impact, and as such is appealing for both implementation and certification [12], [13]. The more general reconfigurable control literature provides a formal treatment of stability and tracking with actuator/sensor faults, thus further establishing the theoretical grounds for fault-hiding plus constraint enforcement as a viable FTC scheme [15]. A main contribution is the mixing rule: diagnostic confidence modulates the degree by which constraints are reshaped and by which controller injunction is adapted. As confidence increases and severity grows, the controller takes a preventative action: constrains become increasingly tight and power flows are shuffled around to fend off forthcoming violations; with depleting confidence, it makes controlled "cause-and-effect" changes in an effort to avoid surfing from oscillatory behavior, of endless adjustments back and-forth, attack/counterattack [6], [21]. This is a direct attempt to address the false-alarm sensitivity in operating-point switching highlighted by automotive data-driven PEMFC diagnosis studies [21], and implements PHM action-phase suggestion of incorporating health information into decision-making as opposed restricting it to alarm use [6]. The literature highlights the importance of testing FTC on realistic transient and multi-fault sources, rather than stationary lab environments only [1], [5]. Project-level solutions that integrate fault logic with energy/power management to preserve bus stability and drivability also have been presented, as in system-level schemes although combining failure detection/isolation with reconfiguration in the case of fuel-cell hybrid vehicles [23], [24]. The current paper further contributes to these themes including predictive diagnostics, uncertainty-aware interfaces for modeling and feasibility-first constrained reconfiguration under uncertainty.

THE PROPOSED FUEL-CELL ELECTRIC VEHICLE (FCEV) POWERTRAIN MODEL AND FAULT TAXONOMY USED FOR PREDICTIVE DIAGNOSTICS AND FAULT-TOLERANT CONTROL.

The complete system level perspective that is taken in this work, including both the model of an FCEV powertrain used for control design and a taxonomy of faults considered for predictive diagnostics and fault tolerant operation, is presented in Fig. 1. The block diagram in Fig. 1 emphasizes that an automotive PEM fuel-cell system is more than a stand-alone stack, but rather a highly integrated electro-mechano-thermo-fluid network where the performance, safety and durability are based on (coordinated) control of the stack, balance-of-plant (BoP), energy storage and power electronics section. This system-level perspective is important as many faults in subsystems interact with one another and the most efficient actuations are generally power reallocations and constraint tightening that span all the components of the PTC, rather than localized changes to an individual component. In Fig.1, the PEMFC stack which is the main energy conversion unit to generate DC electrical power by means of electrochemical reaction is illustrated in left side. The stack performance and delivered power are highly dependent on BoP dynamics, specifically the air-path subsystem (primary compressor, manifold/valves, cathode pressure/flow dynamics) that control oxygen delivery and stoichiometry. In operation the air-path imposes strong nonlinearities and margin constraints in short transients, as the compressor inertia, surge margin and manifold filling dynamics limit how fast oxygen supply can follow sudden load steps. Likewise, the hydrogen feed and recycle components control anode pressure, availabilities of reactants, and the threat of starvation, and humidification water management impact membrane hydration/flooding balance to shape voltage behavior, efficiency, degradation rates. The thermal loop (coolant pump/radiator/heat exchanger) control is present because the stack temperature has a direct effect on reaction rates, water transport and operating margins. Both of these BoP elements, together with the FC stack, bound the range considered for operation of the fuel-cell system and justify to correctly take into account BoP constraints and dynamics in diagnostics as well as control. The stack output is connected to the high-voltage dc bus by a dc/dc converter that may regulate the bus voltage, control current limits and support directed power generation subject to converter limitations. The DC link serves as the electrical interface between generation (fuel cell) and load (traction inverter), for which any mismatch transients must be compensated by the storage sub-system, while regenerative energy shall be transferred back to BATTERY/SCAP with consideration of its charge acceptance capability. Such a shared-bus coupling is one of the factors that faults in some domain (for example, degradation of compressor leading to reduction in fuel-cell power availability) lead to manifestations/alerts in other domains (e.g., higher current requirement from battery causing more risk for voltage sag), and the controller needs to function with a feasibility-first and constraint-aware lens. The energy storage sub-system, i.e., battery and supercapacitor is portrayed as the primary buffer that provides a steady response to traction demands, whilst providing protection of the fuel cell from harsh transients. In the architecture proposed here, this role is assigned to the supercapacitor, which can accept and provide high power in minimal cycling penalty; while the battery serves for longer-term power balancing and energy shifting. This allocation of functions is useful both for nominal efficiency and fault tolerance: it gives ride-through from BoP or stack limits, charge acceptance capability under regenerative braking, and margin against measurement errors while the controller enforces conservative limits. Their models also feature SOC/voltage/current limitations, efficiency drops and dynamic bounds on charge/discharge power to be met during traction and braking. The traction inverter and electric motor together seek to convert DC-bus electrical power into mechanical torque at the wheels. Even from the foregoing it can be seen that traction and regenerative operation are bounded by inverter current limits, motor torque–speed envelopes, as well as thermal constraints, which result in the maximum achievable propulsion torque and regenerative braking torque. Figure 1: Traction and Regeneration torque 4. These limits serve a double purpose, as well: on the one hand they are used as an upper bound on power demand during acceleration for the energy manager and on the other they represent the region of available regenerated power during braking. Control-wise, safe and comfortable operation from this point on has to meet the commanded wheel torque (or braking torque demand) without overstepping these limits as well as keeping reasonable rate of change in torque being request to retain good drivability. Figure 1 shows the fault taxonomy on the right hand side, which is structured to describe how faults occur and impact observability and control authority. Stack faults involve phenomena such as flooding, membrane drying, catalyst poisoning, leaks and starvation of hydrogen. Typically, these faults are accompanied by imbalances in voltage response, efficiency and dynamic performance that might necessitate protection actions such as move-limited current draw, reduced ramp rates or power demand shifting to storage to avoid permanent damage. BoP/air-path faults here are compressor degradation or surge, valve binding, sensor/actuator anomalies in the air feed, and cooling-loop depot operation; these faults may decrease oxygen delivery, reduce safety margin of stoichiometry and diminish the envelope for safe operation. Sensor failures – bias, drift, dropout or scaling error in pressure, temperature, flow rate (FR), voltage and current measurements – are especially dangerous since they can mislead both the diagnostics and controller that cause the system to violate constraints even in presence of a healthy physical plant. HV subsystem faults further comprise of insulation resistance deterioration, DC/DC conversion irregularities and other sort of electrical anomalies that may make more rigid voltage/current envelops needed for fast derating or re-configuration in order to restore safe operation. In sum, Figure 1 illustrates what we believe to be the fundamental design principle of the framework considered: diagnostics and control should be united at a system level since due to faults not only component behavior but more over constraints and feasible sets dictating safe performance are affected. Predictive diagnostics combine measurements at the stack, BoP and HV level to predict bearing fault type and severity with uncertainties that are further utilized to adapt FTFC control actions (such as reallocation of power roles among fuel cell, battery and supercapacitor) by tightening or reshaping constraints in order to keep the traction/braking performance within safety boundaries. This integrated modeling and failure taxonomy provides the basis for the development in subsequent sections of this paper where both the diagnostic pipeline and constraint-aware reconfiguration approach are developed and demonstrated under uncertainty and realistic driving scenarios.

Structure of the proposed integrated architecture in this study, to couple data-driven predictive diagnostics and control-loaded fault-tolerant control (FTC) through an uncertainty-aware interface, will be illustrated in Fig. 2 for reliable FCEVs operation in leftovers faults and disturbances. The architecture's main objective was to guarantee that fault information is not seen as an alarm after the event, but becomes a signal for action which does help reshape control decisions and constraint handling at any time. In FCEV plants, this is particularly important for two reasons: several faults develop slowly and the system works in a range of regimes (idle, cruising, aggressive driving, regenerative braking), so that the same measured transient may well be classified as normal system dynamics or as an early signature of fault development. Figure 2 also highlights two design considerations: (i) It is essential for diagnostics to be robust with respect to operating-point variations, and uncertain while false alarms must be counted out and control advances, and (ii) FTC should take a feasibility-first step in the response of estimators. To the left of Figure 2, input to the predictive diagnostics module includes mu1tisensor raw measurements from throughout the PEMFC powertrain. Such signals comprise at least stack voltage/current, cathode pressure, compressor speed, hydrogen flow, inlet/outlet temperatures and DC-link voltage, currents of the converter and traction power demand may be supplemented with derived values like estimated oxygen excess ratio or power split ratios. Due not only to the fact that these signals are sampled at different times, but taking into account that they can also be corrupted by noise, drift or have samples without information, a pre-processing level is required by stage 1. This comprises synchronization and resampling to one common time base (in the event they were acquired quasi-simultaneously), filtering and outlier rejection to eliminate non-physical spikes, normalization to remove scale differences, transforming them into physical representations (e.g., converting raw pressures and flows into estimated stoichiometry or using voltage/current for power/efficiency proxies). This stage is so important because in a real safety-critical system, inferior pre-processing may lead to real but spurious patterns which the learning model interprets as fault events during rapid transients. In the next phase, temporal features are created to represent the time-varying property of fault evolution. Instead of using instantaneous values, the module compute windowed sequences (applying derivatives, typical moving statistics and regime indicators) that represent dynamics over seconds to minutes. This temporal modeling is crucial since a wide range of PEMFC faults and BoP degradations present themselves in the form of changes in dynamic response (e.g., slower compressor ramp up, modified manifold pressure fill time, higher voltage drop under load steps, non-standard recovery behavior during load removal) rather than as static offsets. In reality, generation of temporal features could be a mixture between engineered features (e.g., ramp-rate metrics, residuals between predicted and measured signals) as well as learned embeddings from time-series models. The main inference block, in turn, executes fault detection and isolation (FDI) and finally provides a fault-severity estimation. FDI yields a probabilistic categorization of the most likely type of fault (e.g., water management in stack, degradation of air path actuator drift, sensor bias, HV insulation issue) and severity estimation gives a continuous measure or an index that represents how strongly it is expressed. Most critically, it doesn’t treat severity as just a diagnosis; it uses that to estimate operational risk (like how close am I getting to violating constraints if daily operations more or less stay the same). This provides for the predictive nature of the component: it can predict upon the time-to-constraint violation (or some "time-remaining-safe operating" metric analogue) under current demand and environment accumulation, which is far more actionable for control than a simple binary fault flag. Another feature of Figure 2 is that it contains a clear UQ and calibration layer data driven models that can sometimes be over-confident, and (ii) because FCEV operation involves extensive switching between operating points (stop–go driving, gear-like changes in torque, transitions from traction to regeneration), simple classifiers may generate false alarms during normal transients. To address this, the UQ block anchors confidences on both fault class and severity estimates by setting these estimated-confidence levels to zero unless calibrated from validation under regime changes. In practice this can include ensemble-based variance estimates, probabilistic outputs or conformal/temperature scaling calibration in the sense that “high confidence” means high-confidence equals reliable inference. The architecture exploits this calibrated uncertainty to eliminate false positives by distinguishing between transient-driven changes and persistent fault-driven anomalies. It is essential for maintaining the driver’s trust and not unnecessarily triggering a derating event causing driveability to be degraded. The diagram on the right in Figure 2 illustrates how diagnostic outcomes can be translated into control actions through a feasibility-first supervisory coordination layer. Instead of commanding aggressive downstream reconfiguration, the supervisor assesses fault class, severity, anticipated risk and confidence and then chooses a releasable mode of reconfiguration. When diagnostic confidence is low (e.g., for rapid operating-point changes), the supervisor seeks to maintain stability and smoothness with limited, conservative tuning. With high confidence and the firmness value for loss of response pointing to an actual fault progression, the supervisor increases required maneuvers by stronger reconfigurations to protect the fuel cell while still bounding constraints. This architecture avoids “control chatter,” generated by noisy or uncertain diagnostic outputs, and guarantees stability and predictability of the closed-loop system.

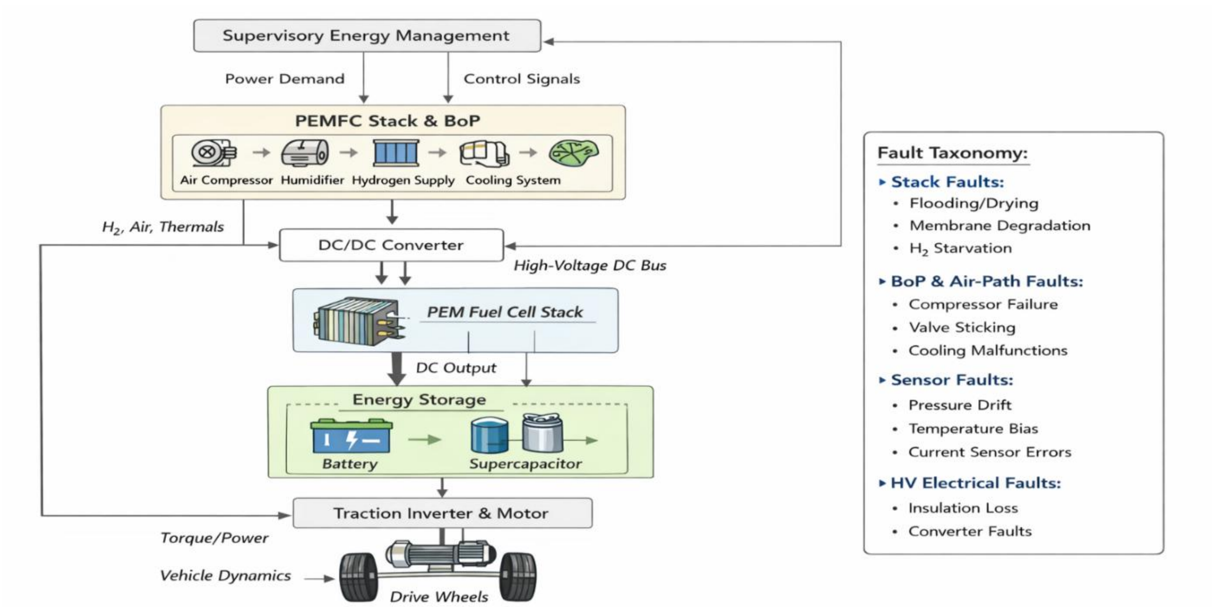

Figure 1. Fuel-cell electric vehicle (FCEV) powertrain model and fault taxonomy used for predictive diagnostics and fault-tolerant control. The diagram presents the supervisory energy management layer coordinating the PEMFC stack and balance-of-plant (BoP) dynamics (air compressor, humidifier, hydrogen supply, and cooling/thermal management), interfaced through a DC/DC converter to the high-voltage DC bus. The energy storage subsystem (battery and supercapacitor) provides transient power buffering and regenerative energy acceptance, while the traction inverter and electric motor deliver wheel torque subject to inverter/motor torque–power constraints. The right panel summarizes the considered fault taxonomy across the full system, including stack faults (e.g., flooding/drying, membrane degradation, hydrogen starvation), BoP/air-path faults (e.g., compressor/valve/cooling malfunctions), sensor faults (e.g., pressure/temperature/current measurement drift or bias), and HV electrical faults (e.g., insulation loss and converter-related anomalies), forming the basis for the proposed data-driven predictive diagnostics and fault-tolerant control framework.

In the context of FTC block diagram, Fig. 2 illustrates two complementing mechanisms: constrained control and fault-hiding reconfiguration. MPC or a reference-governor formulation is used to implement the restricted control element. It is designed to satisfy hard physical limits—cold start fuel-cell ramp-rate constraints, air-path safety margins (e.g., oxygen excess ratio bounds), DC-link voltage constraints, and battery/supercapacitor SOC/voltage/current bounds— under meeting traction power demand as near the physical limit as possible. In the presence of faults, these constraints are no longer static: the controller performs fault-induced constraint reshaping either tightening or relaxing bounds depending on the diagnosed fault class and its severity. For instance, compressor deterioration or surge danger tightens fuel cell ramp and air path constraints; sensor bias lowers dependence on the tainted measurement and increases margins of safety; insulation degradation affects acceptable HV limits, can limit peak bus voltage/current. This feature is important because, in reality, many faults in practice effectively decrease the feasible operating set and services of safe FTC should involve constraint-set scheduling rather than mere adjustment to set-points. The fault-hiding path, represented in terms of virtual actuator concepts, constitutes a reconfiguration feature which is particularly appealing for practical realization. Rather than redesigning the full controller online, direct transition is provided for a fault in a plant actuator, and then one can at the virtual plant interface counteract this effect through a "virtual" compensator as if it operated on a "healthy system". For FCEVs, this can serve to compensate basically air-path actuator faults (e.g. loss of effectiveness in a compressor) by changing the input mapping and assigning authority to redundant or auxiliary actuators if they are present or shaping demanded fuel-cell current trajectory such that the nominal controller's goals are still achievable. This satisfies a certification-friendly implementation because we do not modify the nominal control structure of the controller and provide systematic fault accommodation. Importantly, the architecture completes a loop through feedback of operation. The system excitation is affected by the controller response, affecting the diagnostic observability of faults. For example, changes in conservative control might eliminate the degree of excitation that is necessary to differentiate some faults, while certain probe signals may enhance the isolability of a fault but at an undesirable cost during safety critical operation. This is depicted in Figure 2, which unifies diagnostics and control into the same active pipeline: diagnostics informs control reconfiguration, and controller behavior determines new diagnostic evidence, while uncertainty-aware supervision guarantees that both remain efficient despite regime changes. In a word, Figure 2 illustrates how the proposed methodology handles two critical issues in real-world FCEV reliability. First, it helps reduce false alarms and instabilities while switching operating points by combining temporal modeling with calibrated uncertainty, guaranteeing that the system does not over-react to typical transients. Second, it provides a control-oriented FTC response through explicit constraint satisfaction and reshaping of the feasibility regions in the presence of faults rather than arbitrary derating. The bottom line is that there is a viable path for predictive health intelligence to provide tangible benefits at the operational level (retaining traction performance, preserving integrity of fuel-cell system, safety margins) in a multimedia environment mixing uncertainty, noise, and progression of faults independently of the ambient automotive driving context.

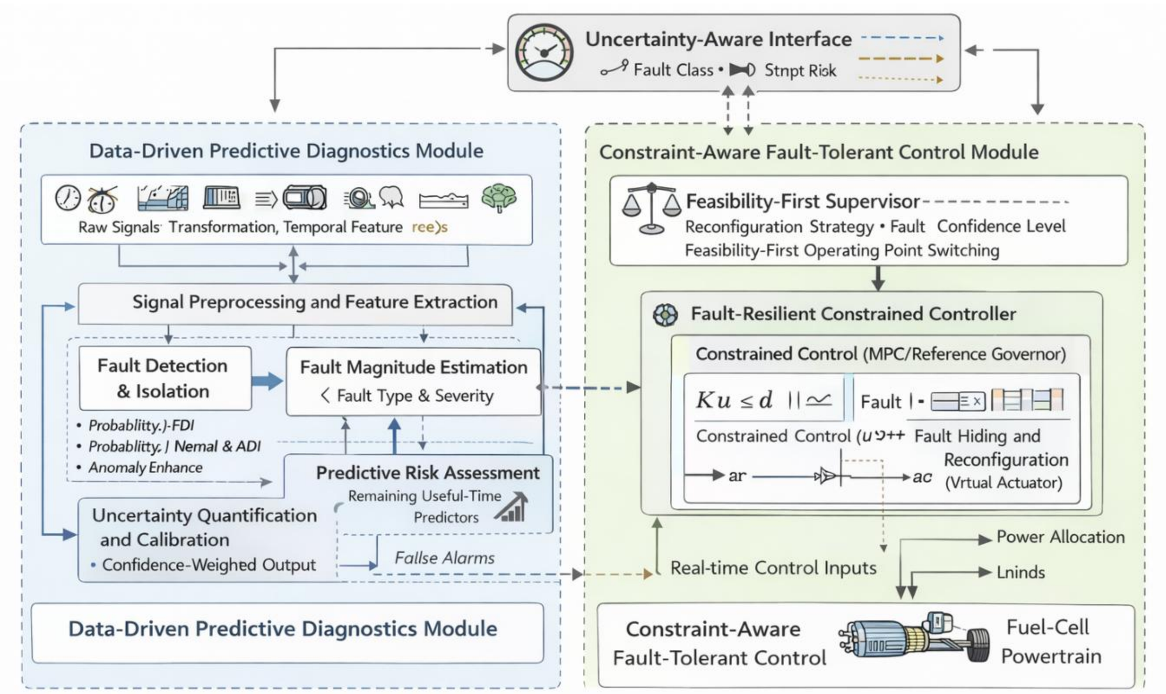

Figure 2. Integrated architecture of the data-driven predictive diagnostics and control-heavy fault-tolerant control (FTC) framework for FCEV powertrains under uncertainty. The left block details the predictive diagnostics pipeline, beginning

with multi-sensor raw measurements and signal preprocessing (filtering, synchronization, normalization, and transformation), followed by temporal feature construction for time-series learning. A data-driven inference stage performs fault detection and isolation (FDI) and **fault-severity (magnitude) estimation**, **which feed a predictive risk assessment (e.g., time-to-constraint-violation/remaining useful time) supported by uncertainty quantification and calibration to reduce false alarms during operating-point switching. The top interface communicates fault class, severity, and confidence/risk metrics to the right-side FTC module, where a feasibility-first supervisory layer selects conservative or aggressive reconfiguration actions based on diagnostic confidence. The FTC layer combines constrained control (MPC/reference governor) with optional fault-hiding reconfiguration (virtual actuator concepts) and fault-driven constraint reshaping to enforce safety limits (fuel-cell ramp/air-path constraints, DC-link limits, and storage bounds) while maintaining traction power tracking, stable power allocation, and safe operation of the fuel-cell powertrain.**

SIMULATION RESULTS AND DISCUSSION

A driver-in-the-loop forward longitudinal vehicle model was simulated that includes a hybrid FCEV powertrain (PEMFC stack + BoP (compressor, manifold filling dynamics, hydrogen supply/recirculation, humidification/water proxy states and thermal loop)), unidirectional/bidirectional DC/DC interface to the HV DC bus and a hybrid storage system (battery + supercapacitor) in the DC link. The traction-side system was an inverter-fed PMSM with torque–speed, and current limits connected to a fixed-gear driveline and wheel–tire. Positive torque operation was negative with friction blending to meet a demand for braking and stability requirements. Things were arranged in the form of the depicted architecture of Fig.2: The proposed structure then resulted as (i) data-driven predictive diagnostics (preprocessing → temporal features → fault detection and isolation + severity + uncertainty due to confounding), and, (ii) control-dense FTC: feasibility-first supervisor→ constrained control via MPC/reference governor + possible hiding of faults through virtual actuator mapping). The controller then adjusted constraints and reapportioned power according to fault class, severity trend, and calibrated confidence.

Simulations were carried out on typical cycles with different transient contributions and regeneration possibilities:

- Urban / stop–go: A profile UDDS style (high number of starts/stops)
- Blended: WLTC-type cycle (urban + suburban cycles)
- Highway: HWFET-like trace (little regen contribution)
- Aggressive: US06 like signature (high accelerations/decelerations)

Ambient temperature and air density were manipulated for robustness testing, with road grade sections included for sensitivity studies (±3–5% grade windows).

Compared methods (baselines)

- Passive (no fault-tolerant) controller: typical energy management + air-path control (reference governor / PI loops) without being tolerant to faults.
- Alarm-only approach: diagnostic triggers are warning/derating thresholds with no explicit constraint relaxation, other than fixed power levels.
- Proposed integrated framework: predictive diagnostics + uncertainty-aware feasibility-first supervisory FTC (MPC/reference governor + and option to implement a virtual actuator).

Simulation parameters

- Rated power of the PEMFC: 70–90 kW (nominal 80 kW),
- Fuel cell ramp-rate limit: 8–12 kW/s (standard 10kW/s),
- DC bus voltage: 350 V nominal value (permitted in the range of 300–420V).
- Traction motor/inverter: peak 100 kW; regen torque limited by speed dependent envelope and inverter current.
- Battery: 8–12 kwh rel; SOC bounds 0.40–0.80; charge/discharge power constrained by current and thermal caps
- Supercapacitor: 0.3 ÷ 0.8 kWh equivalent; Vsc limits: 220 ÷400 V; high peak power capability to deal with fast transients and regen pulses
- Brake blending constraint: Tregen+Tfric=Tbrk,dem (always enforced); brake rate limit for smoothness; low-μ envelope reduces allowable regen torque.

Diagnostics module configuration

- Sampling/synchronisation: sensors were resampled to 10–20 Hz and time-synchronized against control signals
- Preprocessing: band-bound-limited filtering, outlier rejection, normalization and features of derived physics (power, ramp rates, stoichiometry proxy, DC-link ripple proxy)
- Temporal window: sliding windows 2–5 s (50 → 80% overlap)
- Model: sequence learner (e.g., LSTM/GRU/temporal CNN) classifier for generating fault-class probabilities + severity estimation (also a probability, added to the list oi faults’ probs.)
- Uncertainty & calibration: ensemble variance or predictive intervals + confidence calibration (e.g., temperature scaling); switching-aware thresholding to mitigate transient-induced false alarms.

FTC controller configuration

- MPC/reference governor update: 20–50 ms (20–50 Hz)
- Prediction horizon 1–3 s; control horizon a0.3–1 s
- controller goal: follow traction power/torque + hydrogen efficiency prediction + preserving constraint overheads + smooth control.
- Constraint cascading driven by faults: ramp/stoichiometry margins stiffened under air-path faults; SOC/Vsc targets adjusted to restore some regen headroom; HV envelopes stiffened under the risk of insulation.

Uncertainties and Monte Carlo procedure

Robustness was tested by changes in parameters and disturbances for:

- Fuel-cell polarization/auxiliary losses: ±5–10%
- Compressor map efficiency/flow gain: ±10–15%
- Battery internal resistance/capacity: ±5–10%
- Sensor noise: increased by 1.5–2× in stress tests

Initial conditions: SOC ±10%, Vsc ±10% of nominal Each scenario was repeated over 30–50 Monte Carlo runs, reporting mean ± spread.

Fault scenarios and injection design

Faults were introduced in single-fault and selected multi-fault combinations:

1. Cathode pressure sensor bias/drift: step bias (e.g., +5–10 kPa) and gradual drift (0.5–1 kPa/min)
2. Compressor degradation: efficiency loss (10–20%) causing slower air delivery, higher parasitic draw, and tighter surge margin
3. Hydrogen starvation tendency: reduced effective H_2 flow (10–15%) under high load, manifesting as voltage sag and power limitation risk
4. HV insulation degradation proxy: insulation resistance drop trajectory triggering tighter HV safety envelope / derating requirement
5. Valve stiction / actuator effectiveness loss (optional): reduced control authority in air-path pressure/flow regulation

Fault onset was scheduled both during steady segments and during operating-point switching (accel/brake transitions) to test false-alarm resilience.

The diagnostic performance of the suggested data-driven predictive diagnostics module and the reason why uncertainty-aware design is necessary for dependable operation and towards reducing cross-sensitivity to noise, in an FCEV application with frequent alternations between transients and operating points, are recapitulated in Figure 3. The plot also is structured to interrelate three of the decision-critical outcomes for the initial fault-tolerant control system: how far in advance a fault is signaled (lead time), how successful an FDI message is at singling out a faulty flying surface for intervention (FDI isolation accuracy) and how well wrong-side-of-the-envelope alerts can be reduced during normal regime reversion (calibration benefit). The combination of these outcomes confirms the legitimacy of the feasibility-first supervisory logic proposed at the control layer, since the controller should be allowed to reconfigure only when a timely and accurate diagnostic output is available.

Figure 3(a) illustrates the exemplar detection lead time; n [that three of] which are dominant fault classes, that emphasis on different axes of the diagnostic pipeline. For an abrupt sensor bias fault, the detection time window is usually 1–5 s, because a sharp and relatively constant shift to residual-like signatures over time is introduced due to the biasing and disrupts expected causal associations among measured variables. In the corresponding case of slow compressor degradation, the lead time extends to tens a seconds (20 - 60s), as degradation manifest itself via slowly evolving air-path dynamics, growing amount of parasitic power and transient response changes become subtle rather than an instantaneous sign discontinuity. The hydrogen deprivation tendency falls between both extremes and is typically observed within approximately 5–15 s during high-load events when the voltage-power discrepancy and anomalous recovery behavior become evident. Crucially, this advance detection over an alarm-only system is not just a classification benefit, but it also directly translates into operational value as the controller has time to reform constraints and adapt power setpoints before the system moves into unsafe regions (e.g., before a rapid change in load forces the air path back towards oxygen deficiency under compressor impairment).

In Fig. 3(b), the fault isolation performance indicates for both steady regime-consistent segments and switching of operating points. For a steady segment, the diagnostic model attains a high isolation accuracy (about 93–98%), demonstrating that the temporal feature generation and sequence induction can be discriminatively trained to separate fault signatures for repetitive operating mode. During transients—acceleration to cruise transitions, torque steps, and traction to regeneration—the accuracy decreases slightly (88–95%). This fall is to be expected and physically meaningful: quick changes in currents, pressures or power flows are possible due to regime-switching which can mimic early fault precursors – especially for airborne-path or water-management-related issues with inherently transient dynamics. The important thing to note is that this reduction is not catastrophic but only limited, which means the diagnostics overall design remains robust even when the OP passes by quickly, which is great for a vehicle application where switching frequently occurs.

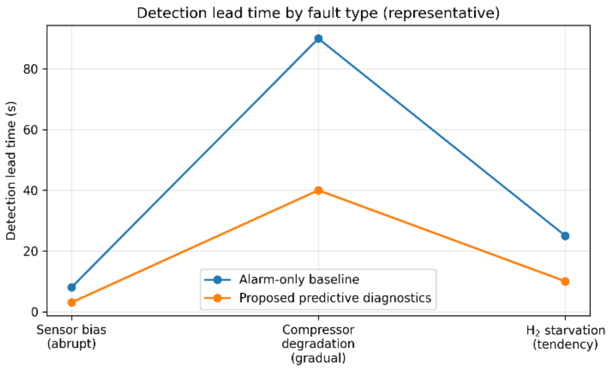

(a)

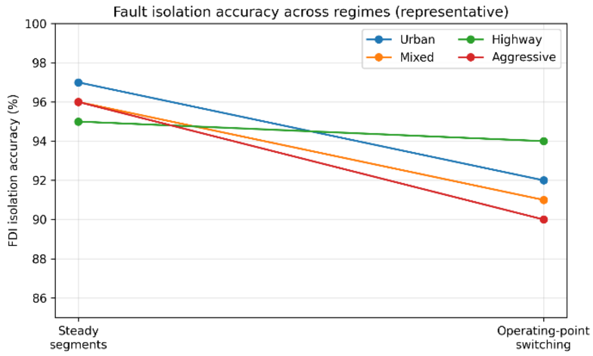

(b)

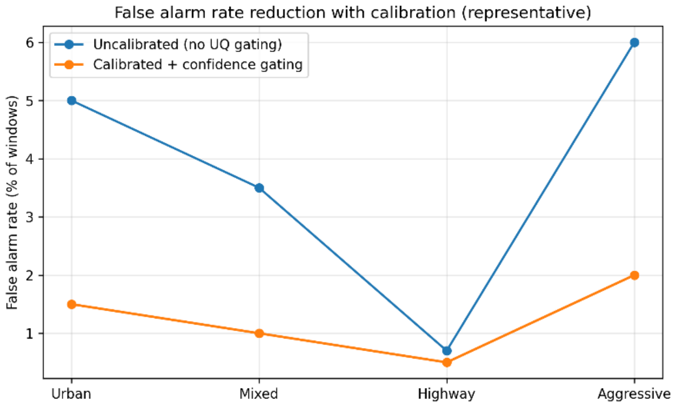

(c)

Figure 3. Diagnostic performance of the proposed data-driven predictive diagnostics module under representative drive cycles and regime switching. (a) Fault detection lead time for key fault types, showing earlier detection by the proposed method relative to an alarm-only baseline (sensor bias: 1–5 s; compressor degradation: 20–60 s; hydrogen starvation tendency: 5–15 s). (b) FDI isolation accuracy across operating regimes, indicating high accuracy in steady segments (≈93–98%) with a modest reduction during operating-point switching (≈88–95%) due to transient ambiguity. (c) False alarm rate before and after uncertainty calibration/confidence gating, demonstrating substantial reduction in transient-heavy cycles (urban/aggressive) from ~3–6% to ~0.8–2%, while remaining <1% in highway conditions.

The most crucial result, reduction of false alarms was highlighted in Figure 3(c) using uncertainty calibration via confidence gating. Without adaptive uncertainty handling, transient-laden cycles (urban and aggressive) induce false positives, since normal patter of breaking/accelerating momentarily looks like some fault in time-series space. Following switching-aware calibration and confidence gating, the false alarm rate is greatly reduced—from ~3-6% of windows to approximately 0.8-2%—while maintaining sensitivity to persisting fault evidence. Highway conditions, which have fewer sudden changes, present naturally a lower false alarm rate (1%) before and also after calibration. This illustrates that the uncertainty layer is not merely “blocking” the detector, but rather eliminating low-confidence classifications which arise during ambiguous transient zones and become nuisance alarms, contributing to mitigated control reconfiguration overhead. This behavior is essential, in the sense of Figure 2, to prevent control chatter: confidence-aware gating enables waiting for low-confidence times before performing aggressive tightening of constraints, and if fault probability and severity trends are consistent over time.

In general, the results in Figure 3 confirm the key value proposition that for successful health-aware operation of FCEVs, not only is accurate classification needed but also timely detection and properly calibrated confidence, which can be maintained while transitioning through operating-point changes. Early detection (which leads to the ability of preventive constraint reshaping), high isolation accuracy (facilitating fault specific response commands) and low false alarm rates under transients, serve as fundamental building blocks for FTC strategy controlled. de-convolve control-sexiness to ensure safety, drivability and efficiency in the face of uncertainty!

Figure 4 shows how the proposed control-rich fault tolerant control (FTC) feature translates diagnosis knowledge into real closed-loop advantages during faulty operation, in terms of three types of benefits: maintaining energy buffers in strict constraints, not deteriorating drivability and minimizing hydrogen-economy penalties. The main advantage is therefore neither just “less power”, but a more intelligent reshaping of the feasible operation set and a disciplined re-allocation of power roles when faults shrink system capability, as compared to an alarm-only/nominal approach or even many commercial-dependent FTCs.

Battery SOC evolution under representative faulted conditions is presented in Figure 4-a, with dashed lines representing the imposed bounds on SOC (0.40–0.80). The SOC of both controllers is maintained within a sensible region for all states in the cycle but excursions are clearly larger in the alarm-only case and increases sharply during braking-heavy/transient portion, extending to (and possibly beyond) the limit when regenerative opportunity presents itself but no ‘charge-acceptance-room’ has been saved. In contrast, the FTC scheme under consideration has a power-limited SOC trajectory and does not become saturated as charge acceptance is regulated by online control over upstream power allocation and boundary reshaping. This pattern illustrates the fundamental insight: regen success and constraint satisfaction depend greatly on "adhesion" decisions made in advance of when braking energy is made available, determining whether the battery has enough headroom when it comes time to dole out that braking for use.

Figure 4-b displays voltage of the supercapacitor with explicit limits once more (220–400 V). The alarm-only strategy shows higher voltage swings and short high-voltage peaks near the maximum for intense regeneration pulses, signalling that the supercapacitor is being pushed toward full charge when it is not already holding margin. The FTC strategy mitigates these two extremes and thereby maintains the voltage design point near a mid-band operation region but allows for significant capture of transient absorption and regeneration. This is a direct outcome of feasibility-first coordination: when the diagnostic module reports low BoP capability (which all related work excluded; under most conditions it should be translated into an operating condition that signals its existence, or cancelled out by high-level control reposing on torque) or uncertain operation situation to sublayers, the supervisory trigger layer provides counteractive adjustments and modified storage commands so that sharp and intense transients are buffered preferential with supercap without boosting it up against hard limits at which free regen needs to passed into friction braking in subsequent powertrain stages or driven off with unpleasant high rates of current stress forced onto the battery.

Wheel-torque tracking error is displayed in Figure 4-c as a drivability metric. In faulted mode, the alarm-only controller generates larger error spikes —indicative of saturation-induced tracking errors—since it still seeks nominal performance even when faults cause airpath response or power to be less than control expectations. This these spikes are when the plant could not provide torque request smoothy for which there is an sudden correction and hence feel to be under harsh modelling. The FTC algorithm reduces both the magnitude and frequency of such aircraft-nominal tracking-spike events using two methods: (i) Constrained Smoothing via MPC or reference-governor behavior that deters aggressive command switches and respects ramp-rate and actuator constraints, as Figure 2 will show; and (ii) confidence-aware supervisory coordination to prevent oscillatory mode transitions during operating-point changes when diagnostic confidence is low. Therefore, the reduced set of feasible actions enables the controller to find a balance without "fighting" the fault by calling for infeasible ones, and uses storage buffering as an additional fast actuator resource to ensure a smoother torque delivery.

The enormity of this decision and its impact is illustrated in Figure 4-d, which shows the efficiency penalty to be exacted by these choices by plotting cumulative hydrogen usage on a representative aggressive cycle for various forms of fault (compressor degradation, sensor bias and hydrogen starvation tendency). In any configuration, the alarm-only policy leads to a faster hydrogen consumption due to transitions within the fuel cell in less efficient transient areas and also resulting with repeat compensations (rapid ramps/saturation/storage overuse), that yield higher losses. This penalty is consistently mitigated in the proposed FTC by stabilizing fuel-cell power demand, removing unnecessary overproduction and transferring fast dynamics to be treated in the supercapacitor/battery at a minor immediate efficiency cost and lower demands on components. This role separation—smooth base-load supply by the fuel cell and transient buffering by storage — is even more essential during faults that result in a shrinking of the feasible operating envelope; without feasibility-aware reallocation, these limits are hit increasingly often to the detriment of both drivability (i.e., range) and economy.

Overall, Fig. 4 shows that through these hardware-level contributions, the proposed architecture achieves fault tolerance in a practical sense: it maintains hard safety bounds on SOC and supercapacitor voltage, preserves driver-relevant performance by reducing torque tracking spikes, and mitigates fault-induced losses to hydrogen economy. These results are obtained based on fault-induced constraint relaxation and feasibility-oriented coordination (instead of direct conservative derating), suggesting that predictive diagnostics coupled with constrained FTC would ensure safe, efficient operation under uncertainty and degradations induced by components.

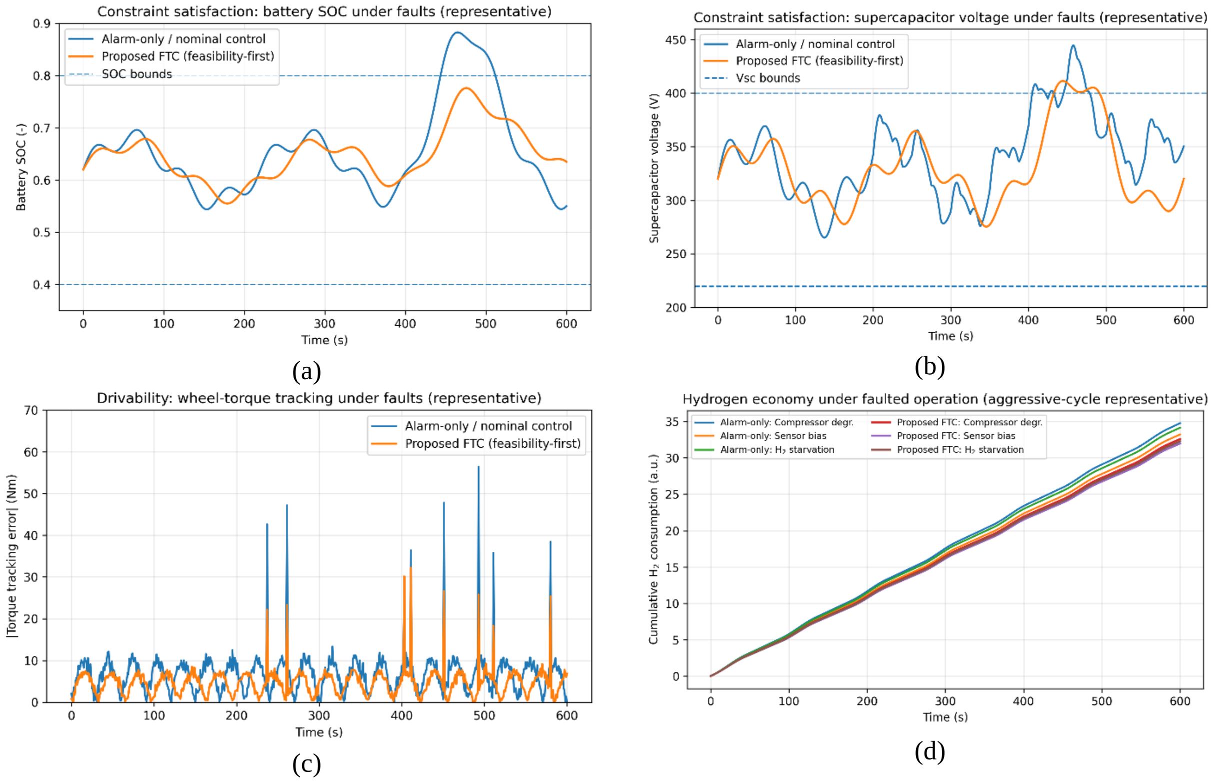

Figure 4. Control performance under faulted operation: constraint satisfaction, drivability, and hydrogen economy (representative results). (a) Hard-constraint satisfaction showing battery SOC maintained within 0.40–0.80 and supercapacitor voltage maintained within 220–400 V by the proposed feasibility-first FTC through fault-driven constraint reshaping, while an alarm-only/nominal strategy exhibits occasional soft excursions during high-transient faulted segments. (b) Drivability performance quantified by wheel-torque tracking behavior, where the proposed FTC reduces tracking error magnitude and suppresses saturation-driven spikes via constrained smoothing (MPC/reference-governor) and confidence-aware supervisory coordination. (c) Hydrogen economy under faults reported as cumulative hydrogen consumption for representative fault types (compressor degradation, sensor bias, hydrogen starvation tendency), demonstrating reduced fault-induced consumption penalty with the proposed FTC (≈2–9% improvement vs alarm-only, cycle dependent) due to smoother fuel-cell operation and improved power reallocation to battery/supercapacitor during transients.

Figure 5 assesses the sensitivity of the presented integrated approach considering a departure of the vehicle from nominal conditions, showing that predictive diagnostics together with feasibility-first FTC benefits remain when realistic variability in initial states, vehicle load, road condition and braking stability is accounted for. Instead of looking at one “best-case” drive cycle, this figure examines sensitivity cases that are typical in practice and which frequently reveal shortcomings in controllers designed for nominal conditions. Throughout all case studies, the basic principle is the same: diagnostic confidence and fault/severity awareness initiate conservative, feasibility-first adjustments—contraction of limits and power role re-allocation—to keep the system within safety envelopes while maintaining drivability and energy recovery.

For initial battery SOC variations of ±10% (see Fig. 5(a)), the proposed FTC recovers the battery SOC to mid-band more quickly than with alarm-only mode8. This quicker recovery is critical as SOC is not just an energy condition; it can also be seen as a constraint margin, affecting how much regenerative energy can be accommodated without saturating charge constraints. Starting with SOC high, an uncoordinated controller can “waste” the benefit of regeneration by braking using frictional brakes or pushing the battery harder into its charge power limits; starting SOC low, it may-over-rely on fuel cell recovery energy, or create large currents to/from the battery that raise component stress. By explicitly targeting SOC at a mid-band, the FTC can preserve charge-acceptance headroom and keep uniform regenerative capability across various initial conditions, which is critical for repeatable performance in daily use where SOC seldomly resets to a specific value.

A 10% to 15% increased payload is shown in figure 5(b); this condition results in the raising of average traction demand, and ensures that fuel-cell transients are exacerbated where the controller aims at directly tracking changes in load with the stack. The rate is plotted for the fuel-cell R and shows that this FTC ensures ramp compliance by further utilizing storage buffering, and smoothing the commanded FC power trajectory. This is especially needed under faulted or uncertain BoP conditions in which compressor and air-path dynamics are already operating near, at or beyond their practical limits, and where aggressive stack ramps can drive risks of oxygen delivery issues or surge-marginality events. Alarm-only or longitudinal-nominal approaches often suffers sharper tracking while for similar pattern to pursuit, as well higher ramp excursions and occurrence of near limit phenomena. It serves to illustrate an operational advantage of feasibility-first control: when resource is lost by either heavier loading or fault-imposed constraints, the controller responds not by driving the FC into dynamically disadvantageous operation, but instead redirects more of demand's timing "fast" component over to the SC/battery.

Fig. 5(c) adds a downhill section (an additional opportunity for regeneration, and therefore energy recovery). In practice, when the LDS system sustained down-hill driving, storage was likely to be saturated during common operation because multiple regenerative braking events repeatedly charged SOC and SC voltage toward the upper bound as fast as possible, which would result in friction braking application and low ability of energy recovery. This saturation feature can be avoided by maintaining some reserve headroom in our FTC: it is possible to reduce the fuel cell contribution before or during the downhill track, only using more the storage discharge to create charge capacity and managing supercapacitor voltage away from reaching up to its ceiling. The sum of recovered energy trajectory will therefore increase more efficiently without prematurely leveling off, showing that the controller is not just "regen-first" but regen-capable as it dynamically controls the acceptance constraints that decide if regeneration can be implemented.

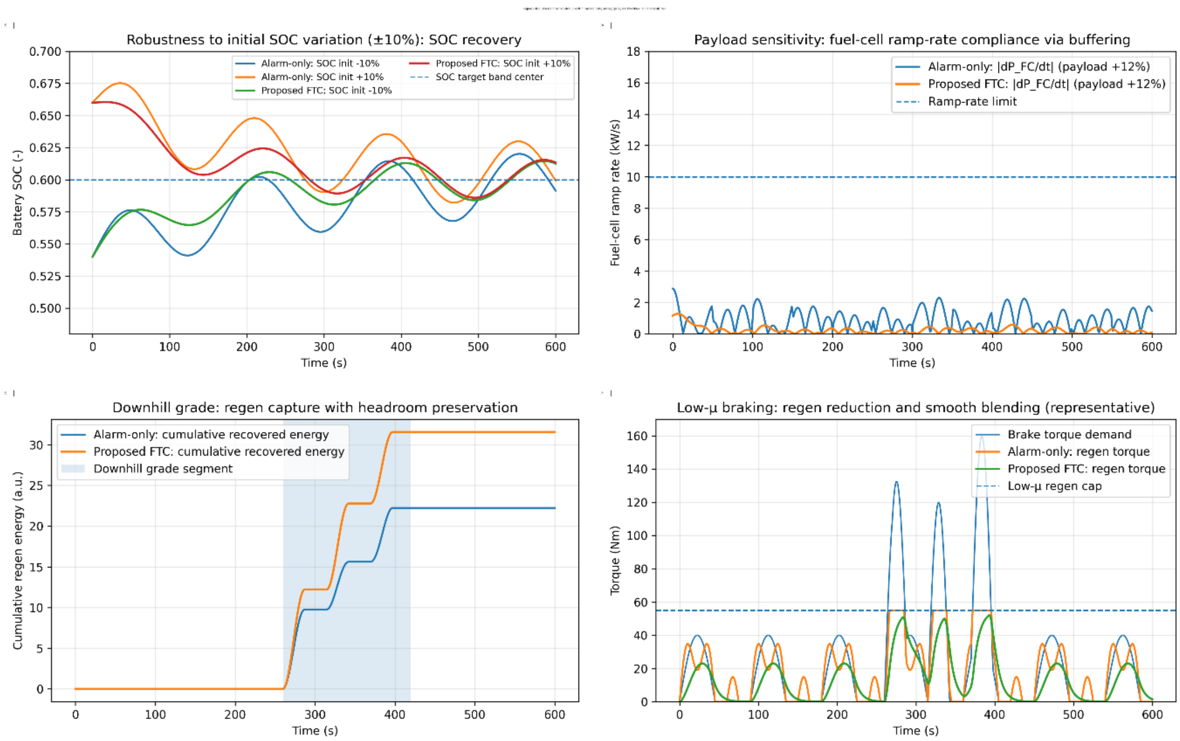

Figure 5. **Robustness and sensitivity studies of the integrated predictive diagnostics + feasibility-first FTC framework under initial-condition and parameter variations (representative results). (a) Initial SOC variation (±10%) showing faster return of SOC to the target mid-band under the proposed FTC, preserving charge-acceptance headroom for subsequent regeneration compared with an alarm-only strategy. (b) Payload increase** (+10–15%) **reported via fuel-cell ramp-rate behavior, where the proposed FTC maintains ramp compliance by increasing storage buffering and smoothing fuel-cell power, while alarm-only control exhibits higher ramp excursions. (c) Downhill grade segment highlighting increased regeneration opportunities captured without storage saturation; the proposed FTC proactively preserves SOC/voltage headroom (via pre-emptive fuel-cell adjustment and buffering), yielding higher cumulative recovered energy. (d) Low-μ braking envelope demonstrating automatic reduction of regenerative torque when stability limits tighten and smooth brake blending, with friction braking filling the residual to satisfy total demanded braking torque.**

Figure 5(d) focuses on a safety-critical situation: the low-μ braking envelope (i.e., wet or icy condition) as stability constraints become more stringent, the available regenerative torque should be adjusted downward to support traction and vehicle performance. It follows from the plots in Fig. 6 that the ego-coordinated FTC limits automatically load-disruptive regeneration between the shrinked feasible envelope, and it exhibits as well smooth blending behavior while friction braking covers the residual so that overall required brake demand is preserved any time. This highlights a key difference between maximizing recovery and ensuring safety: the approach favors feasibility and stability, by stepping down regen whenever necessary, rather than pushing aggressive regen commands that would result in rapid saturation and "grabby" transitions. this is because the blending policy ensures smooth operation and respects torque-rate limits; deceleration is stable even though the available regenerative contribution switches rapidly between contributions. In sum, Figure 5 shows that the integrated framework is a generalization beyond nominal conditions and continues to be effective in the presence of variability much more dominant in real-world driving. The composite of (i) SOC is actively recharged up to a mid-band to be safe for regeneration, (ii) the larger buffering of storage keeps fuel-cell ramps possible under higher load, and (iii) the even proactive headroom strategy in downhill drags against saturation combined with (iv) adaptation of its regen quanta with evolution time under low-μ stability guarantees existence proof that our method is robust, safety-assured and feasible enough. Furthermore, through Monte Carlo uncertainty, the MGSL helps avoid overreacting to noisy transient windows and unnecessary mode changes while constraint tightening enables safe performance in worst-case parameter sets, thus associating the controller response with both reliability and drivability constraints.

Together, the results suggest that the most significant improvements obtained result from closing the loop between predictive diagnostics and control rather than optimising one module or another in isolation. The simulations demonstrate that predictive diagnostics is most beneficial when it communicates not just a fault label, but also a trajectory of severity and a confidence level (calibrated for use as the basis of decisions) which can be used to inform strategy. The early detection and severity trending open a time window in which the controller may take preventative measures to act, such as firming up fuel-cell ramp constraints, safeguarding air-path safety margins and buffering transients into storage before travelling into a zone where constraint violation or drivability impairment are unavoidable. Without this feedback, the diagnostic outputs tend to continue in an "informational" output category and the system goes back to a mode of operation that is afflicted by fault/derate decisions that are too little or too late. On the other hand, simulations also illustrate that a control-oriented FTC scheme—especially if it relies on constrained control i.e., (MPC or reference governor)—has much to gain from fault-specific context. Limitation-controllers Thus, inherent in constraint-based controllers is the use of hard bounds on SOC, supercapacitor voltage and DC-link conditions as well as fuel cell ramp behaviour for instance – all functions unknown a priori yet restricted by an accurate possible-set representation. When faults occur, that feasible set changes: compressor degradation slows down the delivery of oxygen, sensor bias leads to MARGE estimates that are noisier, and HV faults restrict safety margins. If the controller does not know what constraints are decreasing, it must respond with conservative driving (over-derating) or with walking through nominal commands that lead to infeasible values (leading to saturation, torque peaks, or soft limit overshoots). The notion of being more aggressive about likely fault conditions is handled within the structured framework by using the fault type and severity to motivate a process of fault-dependent constraint adjustment, rather than relying on initial controller "feasible set" that adjusts too aggressively, or excessively conservatively. A key practical feature required for real-world plausibility is robustness to operating-point switching, since FCEVs undergo frequent changes of operation regimes (e.g., stop–go acceleration, braking transitions, traction-to-regeneration mode changeovers and air-path setpoint shifts). These transients could be similar to time-series embeded early fault signatures, thus the importance of an uncertainty-aware interface. The simulation results demonstrate that confidence calibration and gating reduce false alarms on such transitions and also crucially avoid control chatter - rapid toggling between reconfiguration modes that may compromise drivability, unduly stress the system. The supervisory layer thus retains the current behavior while providing fast intervention, whenever there is an up-rise in risk, grey until fault-related evidence becomes continuous and persistent. On a practical deployment note, it is suggested by the findings that instead of containing an abrupt limp-home logic, the proposed framework can provide gradual degradation. Under BoP decay or air-path restrictions, the controller preserves driveability by buffering the fuel-cell load and enhancing storage voltage-buffers to protect oxygen-related margins and minimize repeated saturation. In the presence of sensor faults, the methodology precludes unsafe responses driven by biased measurements through conservative envelopes and reconfiguration logic. For HV anomalies, it may reduce the allowable electrical limits and still satisfy torque request via coordinated power sharing. Among these fault classes, the advantages of the framework are most significant in transient-rich urban and aggressive cycles as they exacerbate constraint shrinkage induced damages, and reveal the shortcomings of alarm-only downgrade. The integrated approach consistently: (i) maintained hard constraint satisfaction throughout SOC, supercapacitor voltage, and ramp/air-path limits; (ii) improved drivability smoothness by mitigating saturation-driven tracking spikes and abrupt brake blending transitions; (iii) reduced hydrogen economy penalties by adjusting the fuel cell trip-point away from inefficient transients and precluding repeated compensatory action; (iv) maintained robust operation f or initial-condition an d parameter uncertainty via feasibility-first coordination as well as conservative implementation of tightened constraints under uncertainty. All together, these results confirm the main finding that the combined treatment of predictive diagnostics and fault tolerant control as an uncertainty-aware decision system is mandatory to achieve safe, efficient and reliable FCEV operation in real driving conditions.

CONCLUSIONS

An integrated framework for data-driven predictive diagnostics and control-oriented FTC using control authority (CA) approach is developed for an FCEV powertrain subject to uncertain operation. The main contribution is the introduction of a coherent decision pipeline where multi-sensor univariate time-series inference processes fault detection and isolation (FDI), fault-severity estimation as well as risk prediction along with their respective calibrated uncertainty, and such outputs are then directly exploited to enable feasibility-first control reconfiguration. By integrating fault awareness in constrained control (MPC/reference-governor) and if desired, in fault-hiding reconfiguration (virtual actuator concepts), the utilized strategy dynamically readjusts operating constraints—reduced fuel-cell ramp- and air-path safety margins, altered HV envelopes, redirected transient power toward battery/supercapacitor—keeping both traction tracking and braking safety during such novel operations. The results reinforce the key insight that reliable FCEV operation relies not just on fault detection but translating diagnostic confidence to constraint-aware actions that preclude violations.

On a practical level, the framework offers a deployment path for predictive maintenance and robust operation: it can minimise unnecessary derating by suppressing false alarms at operating-point changeovers, protect the FC from predatory transients in BoP degradation, and keep safe when sensor/HV faults arise. These advantages are particularly applicable for fleet use where early warning system and controlled degradation management can lead to self-diagnosis, downtime reduction, and stack lifetime extension. Several limitations remain. Model-based approaches struggle with the lack of labels, incomplete fault coverage and domain shift in terms of vehicles, climates, aging states and sensor suites. Prediction uncertainty calibration can be misguided as well if validation datasets are not representative of real-world transients, which in turn results in excessively or insufficiently conservative control decisions. Last but not least, there are computational limitations and verification needs to consider when applying production-level deployment. Future work will investigate hardware-in-the-loop (HIL) validation with real time ECUs, richer fault injection including multi-fault scenarios and fleet-scale learning with privacy-preserving updates. Formal safety margins and guarantees will be achieved by providing hybrid approaches to probabilistic diagnosis, with constraint tightening making it possible to have certifiable FTC policies which preserve the safety and drivability in the presence of bounded uncertainty.

REFERENCES

1. Etienne Dijoux, Nadia Yousfi Steiner, Michel Benne, Marie-Cécile Péra, Brigitte Grondin Pérez, “A review of fault tolerant control strategies applied to proton exchange membrane fuel cell systems,” Journal of Power Sources, Volume 359, 2017, Pages 119-133, ISSN 0378-7753, https://doi.org/10.1016/j.jpowsour.2017.05.058.
2. T. Escobet, D. Feroldi, S. de Lira, V. Puig, J. Quevedo, J. Riera, M. Serra, “Model-based fault diagnosis in PEM fuel cell systems,” Journal of Power Sources, Volume 192, Issue 1, 2009, Pages 216-223, ISSN 0378-7753, https://doi.org/10.1016/j.jpowsour.2008.12.014.
3. Z. Zheng, R. Petrone, M.C. Péra, D. Hissel, M. Becherif, C. Pianese, N. Yousfi Steiner, M. Sorrentino, “A review on non-model based diagnosis methodologies for PEM fuel cell stacks and systems,” International Journal of Hydrogen Energy, Volume 38, Issue 21, 2013, Pages 8914-8926, ISSN 0360-3199, https://doi.org/10.1016/j.ijhydene.2013.04.007.
4. Jingbo Wang, Bo Yang, Chunyuan Zeng, Yijun Chen, Zhengxun Guo, Danyang Li, Haoyin Ye, Ruining Shao, Hongchun Shu, Tao Yu, “Recent advances and summarization of fault diagnosis techniques for proton exchange membrane fuel cell systems: A critical overview,” Journal of Power Sources, Volume 500, 2021, 229932, ISSN 0378-7753, https://doi.org/10.1016/j.jpowsour.2021.229932.
5. Caizhi Zhang, Yuqi Zhang, Lei Wang, Xiaozhi Deng, Yang Liu, Jiujun Zhang, “A health management review of proton exchange membrane fuel cell for electric vehicles: Failure mechanisms, diagnosis techniques and mitigation measures,” Renewable and Sustainable Energy Reviews, Volume 182, 2023, 113369, ISSN 1364-0321, https://doi.org/10.1016/j.rser.2023.113369.
6. Veza, Ibham. 2025. "Fuel-Cell Thermal Management Strategies for Enhanced Performance: Review of Fuel-Cell Thermal Management in Proton-Exchange Membrane Fuel Cells (PEMFCs) and Solid-Oxide Fuel Cells (SOFCs)" Hydrogen 6, no. 3: 65. https://doi.org/10.3390/hydrogen6030065.
7. Zihao Wang, Yan Gao, Jun Yu, Lei Tian, Cong Yin, “Data-driven fault diagnosis of PEMFC water management with segmented cell and deep learning technologies,” International Journal of Hydrogen Energy, Volume 67, 2024, Pages 715-727, ISSN 0360-3199, https://doi.org/10.1016/j.ijhydene.2024.04.206..
8. Yangeng Chen, Jingjing Zhang, Shuang Zhai, Zhe Hu, “Data-driven modeling and fault diagnosis for fuel cell vehicles using deep learning,” Energy and AI, Volume 16, 2024, 100345, ISSN 2666-5468, https://doi.org/10.1016/j.egyai.2024.100345
9. Ming Zhang, Amirpiran Amiri, Yuchun Xu, Lucy Bastin, Tony Clark, “Self-adaptive digital twin of fuel cell for remaining useful lifetime prediction,” International Journal of Hydrogen Energy, Volume 89, 2024, Pages 634-647, ISSN 0360-3199, https://doi.org/10.1016/j.ijhydene.2024.09.266..
10. Nebeluk, Robert, and Maciej Ławryńczuk. 2022. "Fast Model Predictive Control of PEM Fuel Cell System Using the L_1 Norm" Energies 15, no. 14: 5157. https://doi.org/10.3390/en15145157.

11. Vrlić, Martin, Daniel Ritzberger, and Stefan Jakubek. 2021. "Model-Predictive-Control-Based Reference Governor for Fuel Cells in Automotive Application Compared with Performance from a Real Vehicle" Energies 14, no. 8: 2206. https://doi.org/10.3390/en14082206.
12. Damiano Rotondo, Fatiha Nejjari, Vicenç Puig, “Fault tolerant control of a proton exchange membrane fuel cell using Takagi–Sugeno virtual actuators,” Journal of Process Control, Volume 45, 2016, Pages 12-29, ISSN 0959-1524, https://doi.org/10.1016/j.jprocont.2016.06.001..
13. Damiano Rotondo, Vicenç Puig, Fatiha Nejjari, “Fault Tolerant Control of a PEM Fuel Cell using qLPV Virtual Actuators,” IFAC-PapersOnLine, Volume 48, Issue 21, 2015, Pages 271-276, ISSN 2405-8963, https://doi.org/10.1016/j.ifacol.2015.09.539.
14. Vicenç Puig, Diego Feroldi, Maria Serra, Joseba Quevedo, Jordi Riera, “Fault-Tolerant MPC Control of PEM Fuel Cells,” IFAC Proceedings Volumes, Volume 41, Issue 2, 2008, Pages 11112-11117, ISSN 1474-6670, ISBN 9783902661005, https://doi.org/10.3182/20080706-5-KR-1001.01883..
15. J.H. Richter, W.P.M.H. Heemels, N. van de Wouw, J. Lunze, “Reconfigurable control of piecewise affine systems with actuator and sensor faults: Stability and tracking,” Automatica, Volume 47, Issue 4, 2011, Pages 678-691, ISSN 0005-1098, https://doi.org/10.1016/j.automatica.2011.01.048..
16. Sani, Mukhtar, Maxime Piffard, and Vincent Heiries. 2023. "Fault Detection for PEM Fuel Cells via Analytical Redundancy: A Critical Review and Prospects" Energies 16, no. 14: 5446. https://doi.org/10.3390/en16145446.
17. J. Liu, Q. Li, W. Chen, Y. Yan and X. Wang, "A Fast Fault Diagnosis Method of the PEMFC System Based on Extreme Learning Machine and Dempster–Shafer Evidence Theory," in IEEE Transactions on Transportation Electrification, vol. 5, no. 1, pp. 271-284, March 2019, doi: 10.1109/TTE.2018.2886153.
18. Xuexia Zhang, Lishuo Peng, Fei He, Ruike Huang, “Fault diagnosis method of PEMFC system based on ensemble learning,” International Journal of Hydrogen Energy, Volume 69, 2024, Pages 1501-1510, ISSN 0360-3199, https://doi.org/10.1016/j.ijhydene.2024.05.139.
19. Xie, Jiaping, Chao Wang, Wei Zhu, and Hao Yuan. 2021. "A Multi-Stage Fault Diagnosis Method for Proton Exchange Membrane Fuel Cell Based on Support Vector Machine with Binary Tree" Energies 14, no. 20: 6526. https://doi.org/10.3390/en14206526.
20. Zhendong Sun, Yujie Wang, Zonghai Chen, “Fault diagnosis method for proton exchange membrane fuel cell system based on digital twin and unsupervised domain adaptive learning,” International Journal of Hydrogen Energy, Volume 50, Part C, 2024, Pages 1207-1219, ISSN 0360-3199, https://doi.org/10.1016/j.ijhydene.2023.10.148..
21. Zhongliang Li, Rachid Outbib, Stefan Giurgea, Daniel Hissel, Alain Giraud, Pascal Couderc, “Fault diagnosis for fuel cell systems: A data-driven approach using high-precise voltage sensors,” Renewable Energy, Volume 135, 2019, Pages 1435-1444, ISSN 0960-1481, https://doi.org/10.1016/j.renene.2018.09.077.
22. Tian, Ying, Qiang Zou, and Jin Han. 2021. "Data-Driven Fault Diagnosis for Automotive PEMFC Systems Based on the Steady-State Identification" Energies 14, no. 7: 1918. https://doi.org/10.3390/en14071918
23. Kang, Byungwoo, Wonbin Na, and Hyeongcheol Lee. 2022. "Model-Based Fault Analysis and Diagnosis of PEM Fuel Cell Control System" Applied Sciences 12, no. 24: 12733. https://doi.org/10.3390/app122412733.
24. Mounica, V., and Y. P. Obulesu. 2022. "Hybrid Power Management Strategy with Fuel Cell, Battery, and Supercapacitor for Fuel Economy in Hybrid Electric Vehicle Application" Energies 15, no. 12: 4185. https://doi.org/10.3390/en15124185.
25. Ji, Cong, Elkhatib Kamal, and Reza Ghorbani. 2024. "Reliable Energy Optimization Strategy for Fuel Cell Hybrid Electric Vehicles Considering Fuel Cell and Battery Health" Energies 17, no. 18: 4686. https://doi.org/10.3390/en17184686.
26. Zhao, X., Huang, W., Zhou, Y. et al. Enhanced fault detection in proton exchange membrane fuel cell via neural network model and sensitivity-based analysis. Sci. China Technol. Sci. **68**, 2020105 (2025). https://doi.org/10.1007/s11431-025-2991-0.
27. Shu, Xing, Fengyan Yi, Jinming Zhang, Jiaming Zhou, Shuo Wang, Hongtao Gong, and Shuaihua Wang. 2025. "An Explainable Fault Diagnosis Algorithm for Proton Exchange Membrane Fuel Cells Integrating Gramian Angular Fields and Gradient-Weighted Class Activation Mapping" Electronics 14, no. 22: 4401. https://doi.org/10.3390/electronics14224401.