Hybrid Fault Detection and Isolation for the Fuel Turbopump Subsystem of the LOX/LNG Expander-Bleed Rocket Engine LUMEN

Eldin Kurudzija
Institute of Space Propulsion, German Aerospace Center (DLR), Köln, Germany

Marcos Quinones-Grueiro
Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, USA

Kai Dresia
Institute of Space Propulsion, German Aerospace Center (DLR), Erlenbach, Germany

Austin Coursey
Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, USA

Gautam Biswas
Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, USA

Jan Deeken
Institute of Space Propulsion, German Aerospace Center (DLR), Hardsthausen, Germany
Institute of Jet Propulsion and Turbomachinery, RWTH Aachen, Germany

Günther Waxenegger-Wilfing
Institute of Space Propulsion, German Aerospace Center (DLR), Hardthausen, Germany
Institute of Computer Science, University of Würzburg, Germany

Abstract

The transition toward reusable and cost-efficient launch systems increases the need for advanced fault detection and isolation (FDI) capabilities to support the reliable and safe operation of liquid rocket engines. Development of these systems typically relies on physics-based simulation models and experimental data. LUMEN (Liquid Upper stage deMonstrator ENgine) is a modular, pump-fed LOX/LNG rocket engine with a nominal thrust of 25 kN, developed and operated by the Institute of Space Propulsion of the German Aerospace Center (DLR). The Rocket Engine Control and Diagnosis Benchmark provides a high-fidelity, experimentally validated EcosimPro/ESPSS simulation model of LUMEN for FDI system development. In this work, we use the fuel turbopump subsystem of this model as a case study and integrate data-driven components into a consistency-based diagnosis (CBD) framework to improve diagnostic performance. The results demonstrate that the hybrid architecture improves isolation accuracy, increasing it from 0.46 to 0.76 compared to the CBD baseline. In addition, we propose a sensor validation network (SVN) that uses a grey-box neural ordinary differential equation surrogate driven by valve-command inputs to decouple sensor faults from component faults. This further increases the isolation accuracy to 0.92, while retaining the ability to detect unknown component faults. These results indicate that incorporating data-driven components helps to compensate for structural non-isolability and improve diagnostic performance while preserving physical interpretability.

2012 ACM Subject Classification Applied computing → Aerospace; Computing methodologies → Anomaly detection; Computing methodologies → Neural networks; Mathematics of computing → Time series analysis

Keywords and phrases Hybrid, Fault Detection and Isolation, Liquid Rocket Engines, LUMEN

Digital Object Identifier 10.4230/OASIcs.DX.2026.8

Acknowledgements ChatGPT (GPT-5.3–5.6) and Claude (Opus 4.6/4.7, Sonnet 4.6/5) were used for language editing and code development. The authors reviewed and verified all generated outputs.

1 Introduction

Liquid-propellant rocket engines (LPREs) are complex, highly coupled systems whose safety depends on reliable monitoring during both steady-state and transient operation. The transition toward reusable and cost-efficient propulsion systems increases the need for advanced fault detection and isolation (FDI) methods to support safe testing, reduce development risk and enable reliable operation throughout the engine life cycle [10, 33, 36]. The Space Shutte Main Engine (SSME) was the first reusable LPRE and incorporated advanced health monitoring capabilities [30]. Beyond conventional redline monitoring, it featured a simulation-based diagnostic algorithm that compared engine measurements against predictions from a nonlinear thermodynamic mass- and energy-balance model to estimate the location and magnitude of predefined anomalies [5]. Such simulation models are developed and validated with experimental data throughout the design, development and testing of rocket engines, providing an established basis for FDI development [25].

More recently, research has expanded beyond this SSME-centered literature to a broader range of LPRE applications. These methods are commonly classified into model-based and data-driven approaches [41]. Quantitative model-based methods have been extensively applied and typically rely on physics-based models for state estimation or parameter estimation [40]. Qualitative model-based approaches, particularly structural-analysis-based methods, have gained increasing attention for exploiting analytical redundancy in the model structure and available sensor configuration. However, their application to LPREs remains challenging due to the large number of candidate residual generators [29]. Moreover, these qualitative approaches characterize fault sensitivity only through residual signatures, without exploiting the distinct physical mechanisms through which sensor and component faults affect system behavior. Leveraging this information could improve isolation performance.

Data-driven FDI methods provide a complementary perspective. They can be divided into unsupervised methods, which typically learn nominal patterns and allow for detecting deviations from such patterns [11, 42], and supervised methods, which assign extracted features to predefined fault classes [31, 39]. Although these methods can learn complex fault patterns, their reliability depends on representative training data and their black-box nature limits physical interpretability. Hybrid FDI methods address these limitations by retaining the model-based structure while selectively introducing data-driven components for residual generation or decision-making [21, 37]. Although such concepts have been investigated in related domains, their systematic application to LPRE FDI remains largely unexplored.

LUMEN (Liquid Upper Stage Demonstrator Engine) is a modular, pump-fed LOX/LNG rocket-engine demonstrator developed by the Institute of Space Propulsion at the German Aerospace Center (DLR). The engine has a nominal thrust of 25 kN and operates at chamber pressures between approximately 40 bar to 80 bar and mixture ratios from 3.0 to 3.8. Hot-fire testing has demonstrated stable operation over a throttling range of approximately 38 bar to 78 bar, establishing LUMEN as a representative experimental test bed for advanced control and FDI research [9]. To support method development before experimental deployment, DLR established the LUMEN Rocket Engine Control and Diagnosis Benchmark which provides an simulation-based environment for designing and evaluating control and FDI methods [8, 32]. It includes an EcosimPro/ESPSS simulation model with a steady-state accuracy of $\pm 5\%$ that enables safe and reproducible evaluation under nominal and faulty operating conditions [23]. In this work, the fuel turbopump (FTP) subsystem of LUMEN is used as a case study. The subsystem-level focus is selected as a first step toward comprehensive engine-level FDI, since turbopump faults constitute a major class of propulsion-system anomalies [4, 14].

The contributions of this work are twofold:

1. We demonstrate how data-driven residual generators and classifiers can be integrated into a qualitative model-based FDI framework to improve diagnostic performance for LPREs.
2. We introduce a sensor validation network (SVN) and combine it with structural analysis to separate sensor faults from component faults, thereby simplifying residual selection while maintaining component-level diagnostic capability.

2 Related Work

The integration of data-driven components into model-based FDI architectures has been investigated in related domains. For fault diagnosis of an internal combustion engine, neural-network-based residuals have been embedded in consistency-based diagnosis frameworks by using structural model information to define grey-box residual generators with physically interpretable inputs, outputs and state dynamics [16, 26]. Recent extensions employ neural ordinary differential equation (neural ODE) models to generate residuals for consistency-based and uncertainty-aware fault diagnosis [18, 27]. A complementary line of work combines model-based residuals with data-driven classifiers in the residual space to improve isolation performance, handle non-ideal residual signatures and support open-set diagnosis [17, 15]. Despite their demonstrated diagnostic value in these domains, such hybrid architectures have not yet been assessed for LPRE.

In LPRE applications, previous work has used Kalman filter residuals to distinguish sensor and actuator from component faults by their different effects on residual statistics [13, 24]. For diagnosis of component faults at the subsystem level, banks of Kalman filters have been employed with different parameterizations representing specific fault hypotheses [13]. Component faults have also been diagnosed using parameter-estimation-based methods, in which component faults are represented by fault factors inferred from deviations between measured engine responses and nominal model predictions [2, 24]. However, these methods for component fault diagnosis require available or reliably reconstructed mass-flow-related quantities at several engine locations, which can be restrictive for flight applications with limited instrumentation due to weight and space constraints [3].

Structural-analysis-based FDI methods incorporate the available sensor configuration directly into the structural model, yielding the detectable and isolable faults for that configuration. For complex LPRE models, however, structural analysis yields a large number of residual-generator candidates, and selection remains nontrivial given the joint constraints of fault sensitivity, causal computability and numerical conditioning [28, 29]. Separating sensor from component faults prior to residual selection removes the need for residual generators to structurally discriminate between sensor and component faults, constraining selection to component faults alone. This motivates combining fault-signature separation with structural analysis to reduce diagnostic ambiguity and identify which component faults remain diagnosable under the available sensor configuration.

3 Case Study

LUMEN, shown schematically in Figure 1, is a modular LOX/LNG expander-bleed engine with two parallel turbopumps [38]. The oxidizer turbopump (OTP) pressurizes LOX and feeds it directly to the main combustion chamber (MCC). Downstream of the LNG pump, the fuel flow splits into two paths: a smaller fraction is routed directly to the MCC and the main flow regeneratively cools the MCC in a counterflow arrangement. The heated coolant

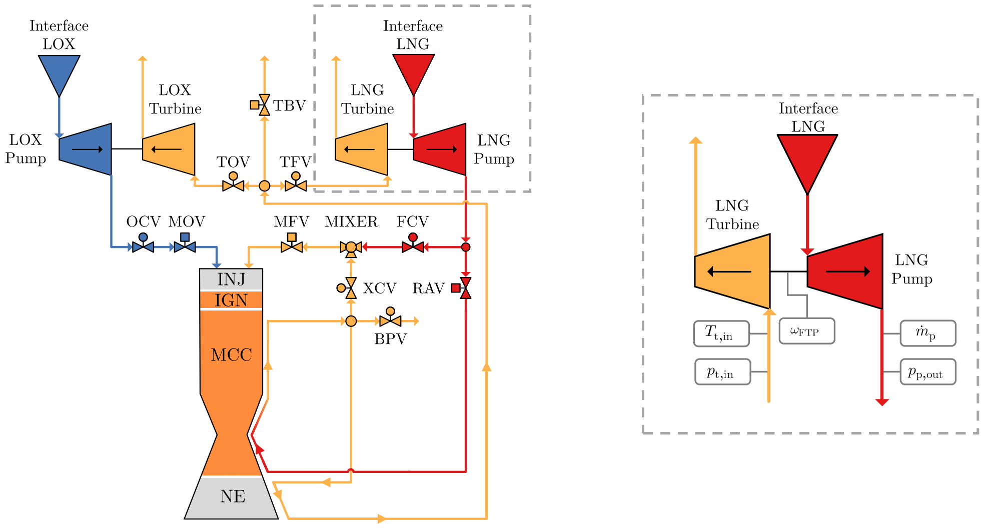

Figure 1 Flow scheme of LUMEN and the considered FTP subsystem with measurements.

is partially remixed with LNG to increase the fuel injection temperature. The remaining coolant regeneratively cools the nozzle extension (NE) before being expanded through the LOX and LNG turbines that drive the respective pumps. The turbine exhaust is then vented without combustion.

At its nominal operating point, LUMEN generates 25 kN of thrust at a combustion-chamber pressure of 60 bar and a mixture ratio of 3.4. In the latest experimental setup, the operating point is set using three electrically actuated control valves: the turbine fuel valve (TFV), the turbine oxidizer valve (TOV) and the fuel control valve (FCV). For preconditioning, ignition and safe shutdown, three on-off valves are used: the main oxidizer valve (MOV), the main fuel valve (MFV) and the regenerative cooling valve (RAV). In future test campaigns, the oxidizer control valve (OCV), mixer control valve (XCV) and bypass valve (BPV) are planned to be reimplemented to enable fault-injection capabilities for validating FDI systems.

In this work, we use the FTP subsystem as the case study. As shown in Figure 1, it comprises the LNG turbine, the mechanically coupled LNG pump and the inlet piping between the LNG interface and the pump inlet. To decouple the FTP from the rest of the engine, we use a sensor configuration consisting of the pump outlet pressure $p_{\mathrm{p,out}}$, pump mass flow rate $\dot{m}_{\mathrm{p}}$, turbopump rotational speed ω_{FTP}, turbine inlet pressure $p_{\mathrm{t,in}}$ and turbine inlet temperature $T_{\mathrm{t,in}}$.

3.1 Mathematical Model

Model-based FDI systems for LPREs typically use a reduced-order physics-based model, in which the engine is represented by interconnected 0-D components governed by first-principles conservation equations for mass, momentum and energy, along with other component-specific relations [1]. The resulting system of ordinary differential equations (ODEs), common in structure across LPREs, captures the engine's nonlinear dynamic behavior once its parameters are identified and calibrated against experimental data obtained during engine development and testing. In the following section, we derive the structural model of the FTP subsystem from the governing equations, following the formulation and simplifications established by Kurudzija et al. [23] through experimental evaluation and modeling of the LUMEN engine.

3.1.1 Turbine

The rotational dynamics of the FTP are described by the angular momentum balance

$$I_{\mathrm{FTP}}\,\frac{d\omega}{dt} = \tau_{\mathrm{T}} - \tau_{\mathrm{P}}, \tag{1}$$

where ω is the rotational speed, I_{FTP} is the FTP inertia and τ_{T} and τ_{P} denote the turbine and pump torques, respectively. The turbine torque is approximated using a performance map that relates the specific torque τ_{specific} to the reduced rotational speed N_{red}

$$\tau_{\mathrm{specific}} = \frac{\tau_{\mathrm{T}}}{r_{\mathrm{t}}\,\dot{m}_{\mathrm{t}}\,a_{\mathrm{in}}} = a + bN_{\mathrm{red}}. \tag{2}$$

The characteristic turbine blade-tip radius is denoted by r_{t}, the turbine mass flow by $\dot{m}_{\mathrm{t}}$, the inlet speed of sound by a_{in} and the polynomial coefficients by a and b. The inlet speed of sound is estimated as a function of temperature T and pressure p, i.e. $a_{\mathrm{in}} = f(T, p)$. The reduced rotational speed is defined by

$$N_{\mathrm{red}} = \frac{r_{\mathrm{t}}\,\omega}{a_{\mathrm{in}}}. \tag{3}$$

Since LUMEN uses supersonic impulse turbines, the turbine mass flow is estimated from the choked-flow conditions at the nozzle throat as

$$\dot{m}_{\mathrm{t}} = \frac{p_{\mathrm{t,in}}\,A_t}{c^*}, \tag{4}$$

where A_t is the throat area and c^* is the characteristic velocity, estimated as $c^* = f(T, p)$.

3.1.2 Pump

The pump performance is described by characteristic curves expressed as functions of the outlet flow coefficient

$$\phi^+ = \frac{\dot{m}_{\mathrm{p}}}{\rho\,\omega}, \tag{5}$$

with $\dot{m}_{\mathrm{p}}$ denoting the pump mass flow. In this work, we assume the fluid density ρ to be constant. The torque coefficient λ^+ is approximated as a linear function of the outlet flow coefficient and rotational speed,

$$\lambda^+ = \frac{\tau_{\mathrm{P}}}{\rho\,\omega^2} = a + b\phi^+ + c\omega, \tag{6}$$

where a, b and c are polynomial coefficients. The pressure rise across the pump is modeled with the momentum balance

$$\left(\frac{L}{A}\right)\frac{d\dot{m}_{\mathrm{p}}}{dt} = p_{\mathrm{in}} - p_{\mathrm{out}} + \rho\,\omega^2\,\psi^+. \tag{7}$$

The term L/A represents the fluid inertia, while p_{in} and p_{out} are the pump inlet and outlet pressures. The head coefficient ψ^+ is modeled as

$$\psi^+ = \frac{p_{\mathrm{out}} - p_{\mathrm{in}}}{\rho_{\mathrm{in}}\,\omega^2} = a + b\phi^+ + c\phi^{+2} + d\omega, \tag{8}$$

where a, b, c and d are polynomial coefficients. The pressure loss in the inlet line between the LNG interface and the pump inlet is described by

$$\left(\frac{L}{A}\right)\frac{d\dot{m}_{\mathrm{p}}}{dt} = p_{\mathrm{if}} - p_{\mathrm{in,p}} - k_{\mathrm{res}}\dot{m}_{\mathrm{p}}^2. \tag{9}$$

The parameters k_{res} and p_{if} denote the piping resistance coefficient and interface pressure, respectively. In this work, we approximate both parameters as constant model parameters.

Table 1 Considered faults and ranges.

Fault	Symbol	Type	Fault Range	Unit
Pump Outlet Pressure Sensor	$f_{p_{\mathrm{p,out}}}$	Sensor	$\pm 5, \pm 10, \pm 15$	bar
Pump Mass Flow Rate Sensor	$f_{\dot{m}_{\mathrm{p}}}$	Sensor	$\pm 0.1, \pm 0.2, \pm 0.3$	kg/s
Turbopump Rotational Speed Sensor	f_{ω}	Sensor	$\pm 2, \pm 4, \pm 6$	1000/min
Turbine Inlet Pressure Sensor	$f_{p_{\mathrm{t,in}}}$	Sensor	$\pm 3, \pm 6, \pm 9$	bar
Degraded Turbine Efficiency	f_{η}	Component	$0.85, 0.90, 0.95$	–
Partially opening BPV (Leakage)	f_{BPV}	Component	$0.4, 0.5, 0.6$	–

3.2 Fault Scenarios

Experimental fault data are typically scarce and costly to obtain for complex safety-critical systems such as LPREs. Simulation-based fault injection is therefore an established, system-agnostic approach to generate representative faulty data across a range of fault types and magnitudes [25]. We represent sensor faults as additive measurement biases, corresponding to a constant offset between the nominal sensor value and the faulty measurement, i.e.

$$y = x + f,$$

where x is the nominal sensor value and f denotes the bias magnitude. The values of the sensor biases are provided in Table 1. For simplicity, we neglect faults in the turbine inlet temperature sensor. This simplification is justified by the scope of the reduced-order FTP model, which focuses on hydraulic and rotational dynamics. The turbine inlet temperature is strongly influenced by upstream heat-transfer processes in the cooling channels, which are not resolved in the model. We introduce a component fault into the FTP subsystem by reducing the torque produced by the turbine

$$\tau_{\mathrm{T}} = f_{\eta}\, \tau_{\mathrm{specific}}\, r_{\mathrm{t}}\, \dot{m}_{\mathrm{t}}\, a_{\mathrm{in}}, \qquad f_{\eta} \in [0.85, 0.95]. \tag{10}$$

Since $f_{\eta} < 1$, this fault represents a degraded turbine efficiency. In addition, we introduce an unknown component fault by partially opening the BPV to simulate a leakage. Due to strong coupling, this fault shifts the engine's operating point. However, since the governing equations of the FTP subsystem remain internally consistent under this fault, the FDI system should be insensitive to it. From the FTP perspective, the subsystem remains consistent with nominal operation. Figure 2 shows the influence of different fault types on various sensors. Although f_{η} increases the mixture ratio, it leads to a reduction in combustion chamber pressure, pump outlet pressure and rotational speed. In contrast, f_{BPV} has only a minor effect on the mixture ratio, but also reduces combustion chamber pressure, pump outlet pressure and rotational speed. Since the engine is controlled in an open loop, a sensor fault affects only the corresponding measurement. In total, we consider six faults, as summarized in Table 1.

4 Methodology

4.1 Structural Analysis

Structural analysis is applied to the mathematical model, defined by the equation–variable incidence relations [12]. Let E denote the set of model equations and Z the set of unknown variables. The structural model is represented as a bipartite graph $G = (E, Z, A)$, with

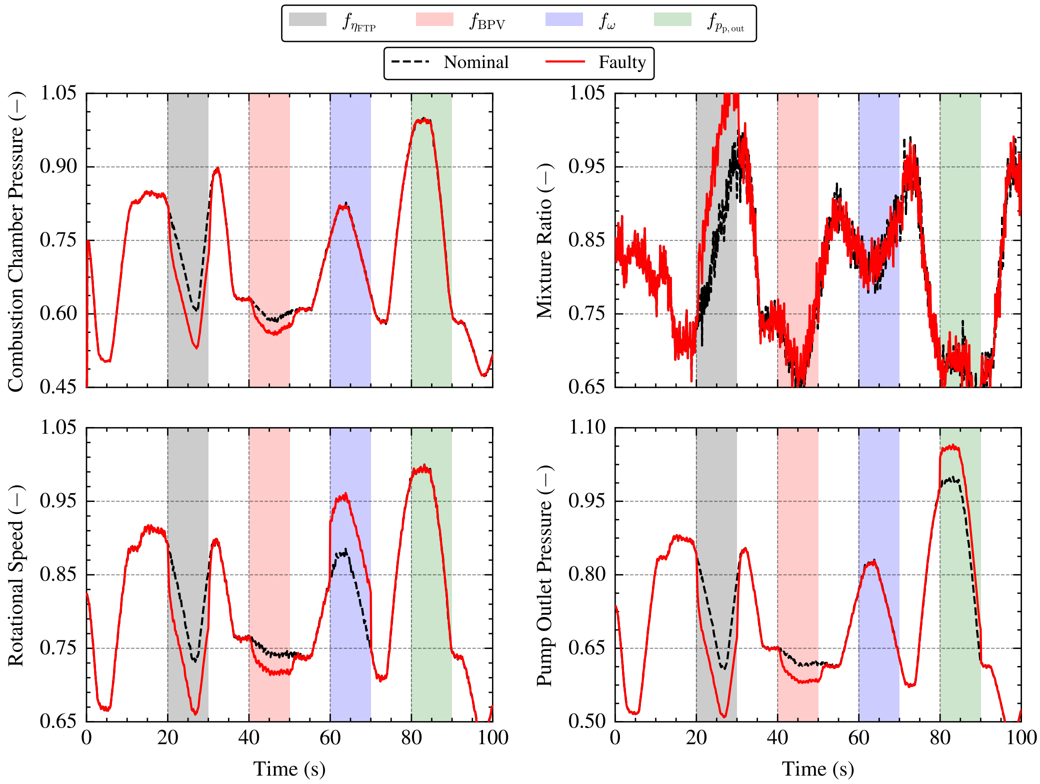

Figure 2 Influence of different fault types on normalized sensor signals. Each signal is normalized by the maximum value of the nominal trajectory.

$(e_i, z_j) \in A$ whenever z_j occurs in e_i. A subset $M \subseteq E$ is structurally overdetermined if M contains more equations than unknown variables, i.e. $|M| > |Z(M)|$, where $Z(M) \subseteq Z$ denotes the unknown variables contained in M. A structurally overdetermined subset M is minimally structurally overdetermined (MSO) if no proper subset $M' \subset M$ is structurally overdetermined [22]. The Dulmage–Mendelsohn decomposition is used to extract the structurally overdetermined part of the model, from which the complete set of MSO sets is derived. These MSO sets constitute residual-generator candidates for structural detectability and isolability analysis. Structural analysis is performed using the Fault Diagnosis Toolbox [12]. For isolability assessment, the considered faults and the measurements are incorporated into the mathematical model. The resulting structural model comprises 18 unknown variables and 20 equations, yielding a redundancy degree of 2 and 5 MSO sets. A simple, minimal hitting set algorithm is then used to select the smallest set of tests that achieves the maximum attainable isolability performance [6]. Figure 3 presents the resulting fault signature and isolability matrix. With four MSO sets, all considered faults are structurally detectable. Since the degraded turbine efficiency fault f_η and the turbine inlet pressure sensor fault $f_{p_{\mathrm{t,in}}}$ produce identical fault signatures in the selected MSO sets, they are structurally indistinguishable.

4.2 Residual Generators

Different methods exist for deriving residual generators from an MSO set. A standard approach is to select one equation of the MSO set as the redundant equation, while the remaining equations form an exactly determined system with an equal number of equations

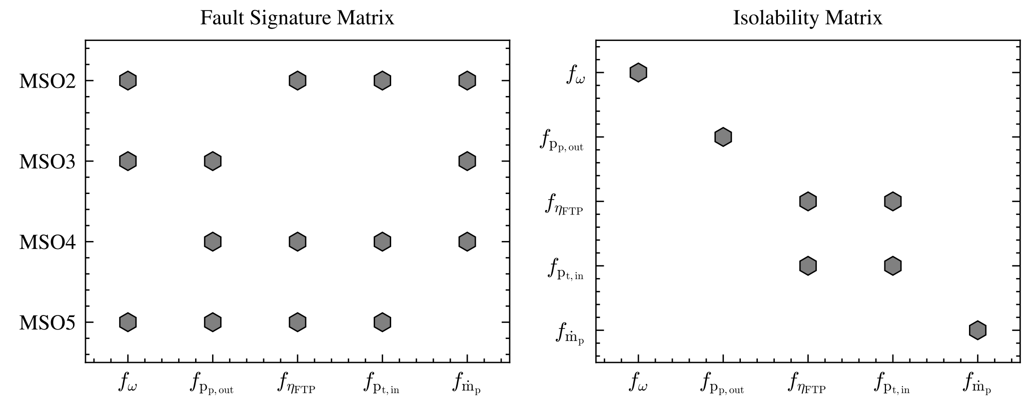

Figure 3 Results of the structural analysis.

Table 2 Derived residual generators and computational causality based on the derived fault signature matrix.

MSO	Redundant Equation	Measurements	Causality	Model
2	y_ω - $\hat{y}_\omega$	$y_{p_{t,in}}, y_{T_{t,in}}, y_{\dot m_p}, y_\omega$	Integral	UKF
3	$y_{p_{p,out}} - \hat{y}_{p_{p,out}}$	$y_{\dot m_p}, y_\omega$	Algebraic	SRG
4	$y_{p_{p,out}} - \hat{y}_{p_{p,out}}$	$y_{p_{t,in}}, y_{T_{t,in}}, y_{\dot m_p}, y_{p_{p,out}}$	Integral	UKF
5	y_ω - $\hat{y}_\omega$	$y_{p_{t,in}}, y_{T_{t,in}}, y_{p_{p,out}}, y_\omega$	Integral	neural ODE

and unknowns. By selecting a measurement equation as the redundant equation, the residual generator compares the model prediction of the measured variable with its observed value. Under nominal conditions, the resulting residual is expected to remain close to zero, i.e. $y - \hat{y} \approx 0$. The computational graph of a residual generator is defined by the sequence in which the unknown variables are computed. Integral causality is obtained if the state variables are computed by integration, while derivative causality is obtained if the state variables are differentiated. If the MSO set contains no dynamic states, the resulting residual generator is algebraic.

In this work, we avoid applying derivative causality to the residual generators, as they require numerical differentiation and are sensitive to measurement noise [20]. Therefore, we select the redundant equation from the MSO sets to obtain residual generators with integral or algebraic causality. Residual generators with integral causality are implemented using a state observer. As the model equations are nonlinear, an unscented Kalman filter (UKF) is employed. UKF-based residual generation avoids explicit model linearization and has previously been investigated for the design of residual generators in LPRE [29]. To enable high fault sensitivity of the UKF residual generators, we choose the process noise covariance matrix to be smaller than the measurement noise covariance matrix.

To design the model-based residual generators, we approximate the mass-flow dynamics (see Eqs. 7 and 9) as quasi-static. The corresponding momentum-balance equations are, therefore, reduced by setting ${}^{d\dot m}/_{dt} = 0$. This approximation is motivated by the characteristic time scale of the momentum-balance dynamics in the mathematical model, which is small relative to the 10 Hz sampling interval of the generated data and, therefore, not resolved. The resulting residual generator for MSO3 has algebraic causality and is implemented as a sequential residual generator (SRG). For MSO5, this approximation yields a computational sequence that requires numerically inverting the nonlinear relationship between flow and the head coefficient (see Eq. 8), which can become ill-conditioned. To avoid these numerical issues, we use a data-driven residual generator as an alternative. For this, we implement a neural ordinary differential equation (neural ODE) residual generator, similar to the architecture in Jung and Westny [18]. The neural ODE residual generator is formulated as a nonlinear state-space model,

$$\begin{aligned}\dot{x}(t) &= \hat{g}(x(t), u(t); \theta_g),\\ \hat{y}(t) &= \hat{h}(x(t); \theta_h),\\ r(t) &= y(t) - \hat{y}(t).\end{aligned} \tag{11}$$

with dynamic state $x(t) \in \mathbb{R}^{n_x}$, inputs $u(t) \in \mathbb{R}^{n_u}$, measured output $y(t) \in \mathbb{R}^{n_y}$, predicted measurement $\hat{y}(t) \in \mathbb{R}^{n_y}$ and residual $r(t) \in \mathbb{R}^{n_y}$. The transition function $\hat{g}$ and the measurement function $\hat{h}$ are parameterized by multilayer perceptrons (MLPs) with trainable parameters θ_g and θ_h, respectively. The state dynamics are integrated over each sampling interval using a fourth-order Runge–Kutta scheme with linear interpolation between samples.

The residual generators used in this work, their redundant equation, input measurements, causality and model type are summarized in Table 2.

4.3 Sensor Validation Network

As discussed in Section 2, sensor and component faults affect the engine in different ways. A sensor fault corrupts the affected measurement, whereas a component fault shifts the engine's operating point, resulting in correlated deviations across multiple sensors. The commanded positions of the three control valves TOV, TFV and FCV set the operating point of LUMEN. For fixed model parameters and in the absence of unmodeled disturbances, the valve-command trajectory defines the engine's nominal behavior and, consequently, the expected sensor outputs at each time step. Based on this observation, we introduce a sensor validation network (SVN) that predicts the nominal sensor output from the control-command inputs. The SVN uses the neural ODE architecture introduced in Section 4.2. The state-transition function $\hat{g}$ propagates the latent state using the valve-command trajectory, while the output map $\hat{h}$ maps the latent state x to the predicted sensor outputs $\hat{y}$.

We perform sensor fault diagnosis by comparing the measured sensor vector with the corresponding SVN prediction. For the n_s considered sensors, the residual vector is defined as

$$\mathbf{r}_{\mathrm{SVN}}(t) = \mathbf{y}(t) - \hat{\mathbf{y}}(t), \qquad \mathbf{r}_{\mathrm{SVN}}(t) \in \mathbb{R}^{n_s}. \tag{12}$$

Since the SVN is driven by control commands rather than by measured signals, sensor faults appear as deviations between $\mathbf{y}$ and $\hat{\mathbf{y}}$.

We then combine the SVN with the structural analysis results to diagnose component faults. Specifically, we use MSO2, which is structurally sensitive to degraded turbine efficiency, together with the SVN-based sensor residuals. This yields distinct signatures for sensor faults, a degraded turbine efficiency as well as unknown component faults. Figure 4 schematically illustrates the resulting diagnostic architecture.

4.4 Fault Detection and Diagnosis

The FDI system operates in two stages on the residuals $\mathbf{r}(t) \in \mathbb{R}^{n_r}$, where n_r denotes the number of residuals derived from structural analysis. The detection stage decides whether the system is in a faulty or nominal operational state. For this, a per-residual two-sided

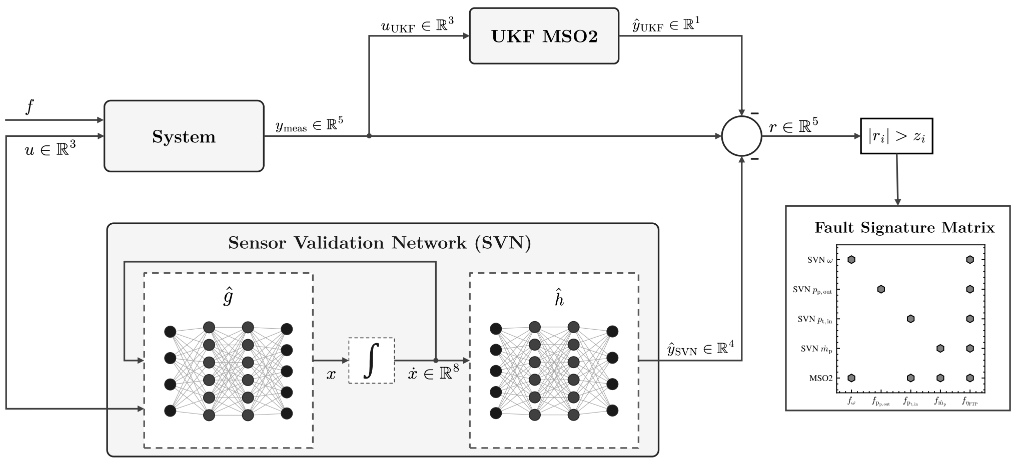

Figure 4 Schematic overview of the SVN-based diagnosis architecture.

hypothesis test is applied to the standardized residuals, flagging residual i when

$$\frac{|r_i(t) - \mu_{0,i}|}{\sigma_{0,i}} > z_{1-\alpha/2}, \tag{13}$$

where $\mu_{0,i}$ and $\sigma_{0,i}$ denote the mean and standard deviation of residual i under nominal operation and $z_{1-\alpha/2}$ the $(1-\alpha/2)$-quantile of the standard normal distribution. A fault is detected when at least one residual is triggered, i.e. $d(t) = \mathbf{1}[C(t) \neq \emptyset]$. This provides a set of triggered residuals $C(t) \subseteq \{r_1, \ldots, r_{n_r}\}$, which are evaluated in the next stage.

If a fault is detected, fault isolation is performed in the second stage to compute the set of diagnosis candidates. For isolation, we evaluate two different approaches. The first follows a consistency-based fault isolation logic (CBD) [7, 34]. Following the structural derivation of the residuals, the single-fault setting and a calibrated detection stage, the exoneration and no-false assumption are justified. Therefore, f_k is admitted as a diagnosis candidate iff $C(t) = S_k$, where $S_k \subseteq \{r_1, \ldots, r_{n_r}\}$ denotes the set of residuals expected to trigger under fault f_k, yielding the candidate set

$$\mathcal{D}(t) = \{f_k \in \mathcal{F} \mid C(t) = S_k\}\,. \tag{14}$$

The limitation of this logic for the FTP subsystem is that the degraded turbine efficiency fault f_η is structurally not isolable from the turbine inlet pressure sensor fault $f_{p_{\text{t,in}}}$. To use information beyond the binary structural fault signature, we use a data-driven isolation logic based on a set of fault-specific one-class support vector machines (OC-SVMs), similar to Jung et al. [17]. OC-SVMs are trained on samples from a single class only and and learn a decision boundary around this class in the feature space, classifying new observations outside the boundary as novel or anomalous [35]. For each fault $f_k \in \mathcal{F}$, an OC-SVM is trained on residual data generated under that fault, learning a closed acceptance region that encloses the corresponding residual vectors and thereby models a single fault hypothesis. A fault is selected as a diagnosis candidate if the residual vector lies in the acceptance region of the corresponding classifier, i.e.

$$\mathcal{D}(t) = \{f_k \in \mathcal{F} \mid s_k(\mathbf{r}(t)) \geq 0\}\,, \tag{15}$$

where $s_k(\mathbf{r})$ denotes the decision score of the classifier trained for fault f_k and $s_k(\mathbf{r}) \geq 0$ indicates that the residual vector is consistent with the corresponding fault hypothesis. If the fault-specific distributions define separable regions in the residual space, the corresponding fault modes can be isolated as distinct candidates. If these regions overlap, the affected fault hypotheses remain jointly in the candidate set $\mathcal{D}(t)$. An empty candidate set yields an *unknown* diagnosis, i.e. a fault is detected in the first stage, but no candidate is accepted in the isolation stage.

To reduce the influence of transient effects on the diagnosis, we stabilize both detection and isolation decisions by M-of-N persistence filters [15]. These filters aggregate the most recent decisions over a sliding window of length N, accepting a decision as persistent only if it occurred in at least M of these samples [19].

4.5 Evaluation Metrics

We evaluate the FDI system using four metrics, two for the detection stage and two for the end-to-end isolation. Let $\mathcal{N}$ denote the set of samples generated under nominal operation and $\mathcal{F} = \bigcup_k \mathcal{F}_k$ the set of samples generated under a known fault, where $\mathcal{F}_k$ contains the samples corresponding to fault f_k. The detection stage is evaluated based on the false alarm rate (FAR) and the true detection rate (TDR),

$$\mathrm{FAR} = \frac{1}{|\mathcal{N}|} \sum_{t \in \mathcal{N}} \mathbf{1}[d(t) = 1], \quad \mathrm{TDR} = \frac{1}{|\mathcal{F}|} \sum_{t \in \mathcal{F}} \mathbf{1}[d(t) = 1], \tag{16}$$

where $d(t) \in \{0, 1\}$ is the detection indicator. The end-to-end isolation performance is evaluated based on the isolation accuracy IA and its top-k relaxation IA_k,

$$\mathrm{IA} = \frac{1}{|\mathcal{F}|} \sum_{t \in \mathcal{F}} \mathbf{1}[\hat{c}(t) = c(t)], \quad \mathrm{IA}_k = \frac{1}{|\mathcal{F}|} \sum_{t \in \mathcal{F}} \mathbf{1}[c(t) \in \mathcal{D}(t) \wedge |\mathcal{D}(t)| \leq k], \tag{17}$$

where $\hat{c}(t)$ is the final diagnosis label, $\mathcal{D}(t)$ the diagnosis candidate set and $c(t)$ the ground-truth fault class. In this work we use IA_2, i.e. accepting an ambiguity group of up to two candidates as a successful diagnosis.

4.6 Data Generation

The data generation process aims to produce representative simulated trajectories comparable to those used during ground tests while providing denser coverage of the operational domain. For this, we split the process into two stages. In the first stage, steady-state data are uniformly generated within LUMEN's operational envelope using Latin Hypercube Sampling (LHS). To set the sampled operating point, two PI controllers are designed for TOV and TFV. The position of FCV is set as a function of the operating point in the open loop. In the second stage, transient trajectories are generated by connecting randomly selected steady-state operating points with ramps with randomized transition times. The resulting load point is held for a randomly sampled dwell time. To increase data diversity, the FCV position is randomly perturbed within a $\pm 5\%$ range. Each generated trajectory has a duration of 100 s. For data generation, the simulation model is initialized from a steady-state solution and simulated in open loop using sampled valve positions with a timestep of 0.1 s.

For training and validation of the data-driven isolation stage, we generate 30 trajectories for each fault class and for the nominal class in each data split. The faults are introduced at $t = 20$ s as steps and remain active for the rest of the trajectory, simplifying fault labeling and evaluation.

4.7 Implementation Details

Both the neural ODE residual generator and SVN are trained on 250 fault-free trajectories. Each trajectory is segmented into overlapping windows of approximately 50 s length with 50 % overlap. All inputs and outputs are standardized to have zero mean and unit variance, based on the training set statistics. Following Mohammadi et al. [27], we use a curriculum-learning strategy, where the trajectory-window length is increased from 20 s to approximately 50 s using a piecewise-linear schedule. Model parameters are optimized using Adam to minimize mean squared error, with an initial learning rate of 10^{-3} that decays exponentially by a factor of 0.99 per epoch. Early stopping is applied based on the validation loss.

4.7.1 State Initialization

For the neural ODE residual generator, we use a noise-aware initialization, similar to that of Mohammadi et al. [27]. During training, the initial dynamic state is sampled as $x_0 \sim \mathcal{N}\left(y_0, \sigma^2\right)$, centered at the first measurement of each trajectory window with σ the measurement-noise standard deviation. During inference, the first measurement initializes the state directly. For SVN, we use a different initialization strategy. Assuming that the simulation starts from a steady-state solution, the initial state is fixed to $x_{t0} = 0$. This assumption is justified by the results in Kurudzija et al. [23], which demonstrate that the engine dynamics can be separated into slow thermal transients and faster fluid dynamics. As the simulation model is initialized from a thermal steady-state solution (see Section 4.6), the surrogate model is trained to capture the system's fast dynamics, which reach a steady-state solution rapidly after changes in the valve commands. To avoid penalizing initialization transients, the loss is evaluated after an initial 2 s warm-up period within each trajectory window.

4.7.2 Hyperparameters

The architecture and training hyperparameters of both neural-ODE-based models are selected empirically. The selection is based on validation set performance using 50 fault-free trajectories, with the objective of balancing training time and long-horizon rollout accuracy. The selected MLP architecture consists of two hidden layers, each with 128 neurons. For SVN, the latent-state dimension is set to 8. Following Jung and Westny [18], SiLU activations are used in $\hat{g}$, while $\hat{h}$ uses ReLU activations with layer normalization after each nonlinear layer. The hyperparameters of the detection and isolation stages are selected by a grid search on the validation set consisting of 30 trajectories for each fault and the nominal class. The detection threshold is chosen to satisfy FAR ≤ 0.02 while maximizing TDR. For the set of OC-SVMs, we use the RBF kernel and select its hyperparameters ν and γ to maximize the per-class isolation accuracy. The parameter ν upper-bounds the fraction of training samples allowed to lie outside the decision boundary, controlling the trade-off between boundary tightness and tolerance to training outliers, while γ determines how far the influence of an individual training sample extends, with larger values producing a more localized and flexible decision boundary. To ensure a consistent evaluation of the isolation stage, we apply the same persistence filter parameters to all architectures, with $M = 6$ and $N = 10$.

5 Results and Discussion

We evaluate the proposed FDI architectures on a test set of 30 trajectories per fault class and 30 nominal trajectories, using the metrics introduced in Section 4.5 to quantify detection and end-to-end isolation performance. We further assess performance on the unknown fault

Table 3 Detection and isolation performance of the evaluated FDI systems. The best value for each metric is highlighted in bold.

System	FAR	TDR	IA	IA_2
Hybrid	**0.00**	0.96	0.76	**0.94**
CBD	**0.00**	0.96	0.46	0.85
SVN	0.01	**0.99**	**0.92**	0.92

class and analyze the time-resolved diagnostic output along an exemplary trajectory. The following abbreviations denote the evaluated FDI systems:

- Hybrid: fault isolation using an OC-SVM per fault hypothesis, applied to MSO-derived residuals.
- CBD: fault isolation via consistency-based diagnosis, applied to MSO-derived residuals.
- SVN: fault isolation via consistency-based diagnosis, applied to residuals from the proposed sensor validation network combined with MSO2.

For the SVN architecture, the data-driven isolation stage is omitted, as the SVN residuals render the considered fault classes structurally isolable.

5.1 Detection and Isolation Performance

Table 3 summarizes the detection and isolation performance of the FDI systems. While the detector operating on MSO residuals achieves a perfect false alarm rate (FAR = 0.00), its true detection rate is slightly lower than that of the SVN-based detector, i.e. TDR = 0.96 compared to 0.99. SVN achieves a FAR of 0.01, indicating that both detectors perform comparably overall. The isolation performance varies substantially across the three architectures. As expected, the CBD baseline achieves the lowest isolation accuracy (IA = 0.46) as f_η and $f_{p_{\mathrm{t,in}}}$ are not structurally isolable. Hybrid yields a substantial improvement to IA = 0.76. This indicates that the residual distributions of f_η and $f_{p_{\mathrm{t,in}}}$ are sufficiently separated for data-driven isolation. SVN achieves the highest isolation accuracy at IA = 0.92. Allowing an ambiguity group of two candidates partly closes the gap for Hybrid and CBD. Hybrid achieves the highest isolation accuracy ($\mathrm{IA}_2 = 0.95$), followed by SVN ($\mathrm{IA}_2 = 0.92$) and CBD ($\mathrm{IA}_2 = 0.85$). For CBD, the gap between IA = 0.46 and $\mathrm{IA}_2 = 0.85$ reflects the structural indistinguishability of f_η and $f_{p_{\mathrm{t,in}}}$. For SVN, $\mathrm{IA} = \mathrm{IA}_2 = 0.92$ shows that the candidate set is consistently a singleton, with more discriminative residual generation enabling unique isolation. For Hybrid, the increase from IA = 0.76 to $\mathrm{IA}_2 = 0.94$ suggests that the data-driven isolator resolves the structural indistinguishability for most samples but leaves an ambiguity in a small fraction of cases, where the distributions of f_η and $f_{p_{\mathrm{t,in}}}$ overlap in the residual subspace and both faults remain candidates.

This is illustrated in Figure 5. The residual response of $f_{p_{\mathrm{t,in}}}$ exhibits two directions depending on the fault sign, i.e. a positive and a negative bias. On the other hand, the residual response of f_η is unidirectional. The overlap between the two fault classes in the residual space occurs specifically when $f_{p_{\mathrm{t,in}}}$ exhibits a sensor fault with positive bias, as highlighted in Table 4. Hybrid achieves correct and unambiguous isolation whenever $f_{p_{\mathrm{t,in}}}$ has a negative bias, but fails to do so when it has a positive bias. However, the isolation performance increases with fault magnitude in the positive-bias region, as larger magnitudes shift the residual vector further into the separable region of the residual space. Specifically, the isolation accuracy increases from IA = 0.09 at +3 bar to IA = 0.95 at +9 bar. For

Table 4 Fault-magnitude-wise isolation performance (IA) for $f_{p_{\mathrm{t,in}}}$ and f_η. The best value for each fault magnitude is highlighted in bold. The CBD architecture is not shown as these faults are not structurally isolable.

	$\mathbf{f_{p_{t,in}}}$ **(bar)**						$\mathbf{f}_\eta(-)$		
Model	+3	−3	+6	−6	+9	−9	0.95	0.90	0.85
Hybrid	0.09	0.95	0.57	0.95	0.95	0.93	0.52	0.07	0.38
SVN	**0.99**	**0.99**	**0.99**	**0.99**	**0.99**	**0.99**	**0.97**	**0.99**	**0.99**

f_η, the per-magnitude isolation accuracy exhibits a less consistent pattern, with IA $= 0.07$ at $f_\eta = 0.90$, IA $= 0.52$ at $f_\eta = 0.95$, and IA $= 0.38$ at $f_\eta = 0.85$. In contrast, SVN achieves correct and unambiguous isolation across all magnitudes and fault classes with near-perfect accuracy (IA $= 0.99$). This is consistent with the design of the SVN architecture and indicates that incorporating fault-discriminative information into the residual generation stage substantially improves isolation performance for structurally similar faults.

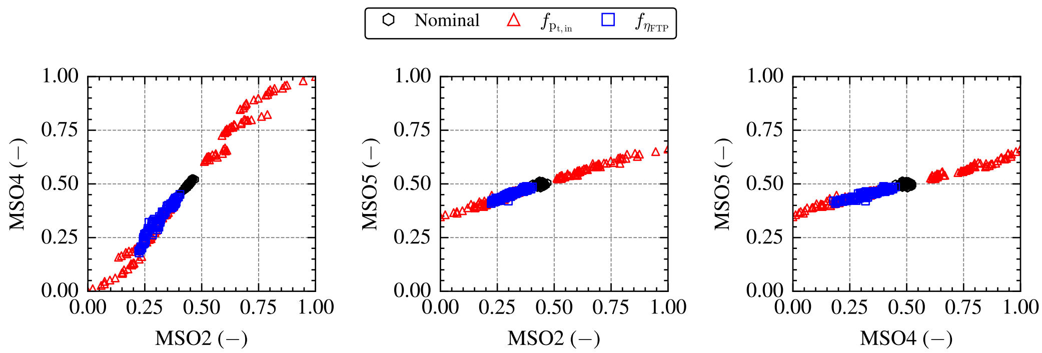

Figure 5 Response of $f_{p_{\mathrm{t,in}}}$ and f_η in the residual subspace. For better readability, the residuals are normalized and only every 100−th point is illustrated.

Table 5 shows the per-fault isolation performance. As discussed, the performance gap between the three configurations is most pronounced for $f_{p_{\mathrm{t,in}}}$ and f_η. For f_ω and $f_{p_{\mathrm{p,out}}}$, all three configurations achieve near-perfect isolation accuracy (IA ≈ 0.99), indicating that these faults are well-separated. For $f_{\dot{m}_{\mathrm{p}}}$, the performance of the three configurations differs. Hybrid achieves the highest isolation accuracy (IA $= 0.83$), while CBD and SVN reach only 0.34 and 0.63, respectively. This difference can be attributed to the sensitivity of the isolation mechanism to fault magnitude, as illustrated in Figure 6. At small magnitudes, $f_{\dot{m}_{\mathrm{p}}}$ does not produce a sufficiently strong response in the residuals to trigger the required crossing pattern and, therefore, fails to admit $f_{\dot{m}_{\mathrm{p}}}$ as a diagnosis candidate. As a result, both architectures provide a *unknown* diagnosis. In contrast, the hybrid architecture operates directly on the residual vector $\mathbf{r}(t)$. The per-fault classifiers were trained on $f_{\dot{m}_{\mathrm{p}}}$ data across all magnitudes and can therefore identify the fault even when the residual response is weak. This illustrates a key advantage of data-driven isolation: by decoupling fault isolation from the detection threshold, the per-fault classifier retains sensitivity to low-magnitude faults that would otherwise go undetected by threshold-based crossing logic.

Table 5 Per-fault isolation performance. The best value for each fault class is highlighted in bold.

System	IA f_ω	$f_{p_{\mathrm{t,in}}}$	$f_{p_{\mathrm{p,out}}}$	f_η	$f_{\dot{m}_\mathrm{p}}$
Hybrid	0.96	0.74	0.97	0.32	**0.83**
CBD	**0.99**	0.00	**0.99**	0.00	0.34
SVN	**0.99**	**0.99**	**0.99**	**0.99**	0.63

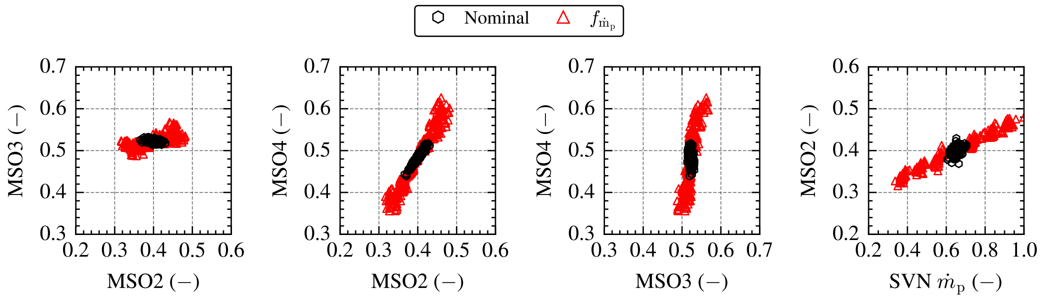

Figure 6 Response of $f_{\dot{m}_\mathrm{p}}$ in the residual subspace. For better readability, the residuals are normalized and only every 100−th point is illustrated.

5.2 Unknown Fault Performance

Table 6 reports the performance of all three configurations for the unknown fault class f_BPV across different fault magnitudes. The fault magnitude $f_\mathrm{BPV} \in [0.4, 0.6]$ denotes the relative opening of BPV. The architectures exhibit different behaviors for the unknown fault class. Since the derived MSOs for Hybrid and CBD insensitive to the unknown fault class, the detection stage should not be triggered, as the FTP is operating nominally. Both architectures effectively reject the unknown fault class during the detection stage, resulting in zero false alarms ($\mathrm{FAR}_\mathrm{unk} = 0.00$). This confirms that the MSO-based architectures do not misclassify the unknown fault class. Further, this implies that the fault is undetected within the FDI system for this subsystem, which is desirable for a distributed diagnosis system where this fault class is handled by another local diagnosis system.

In contrast, SVN is sensitive to this fault class, as it produces a correlated deviation across multiple sensor signals, captured by the SVN residuals. Since MSO2 is structurally insensitive to this fault, the isolation logic under the exoneration and no-false-alarm assumptions produces an empty candidate set, which the pipeline resolves as an *unknown* diagnosis. SVN achieves $\mathrm{IA}_\mathrm{unk} = 0.98$, confirming reliable detection of the unknown fault class. The marginal shortfall is attributable to the warm-up phase of the persistence filter, during which the first $N - 1$ samples of each fault trajectory do not yet contribute a fully populated sliding window, slightly reducing the per-sample rejection rate below unity.

5.3 Trajectory-Level Performance

To evaluate the trajectory-level performance, the exemplary test run introduced in Section 3.2 is used. The run consists of four sequential fault episodes: $f_\eta = 0.9$ is active during $t \in [20, 30]$ s, $f_\mathrm{BPV} = 0.5$ during $t \in [40, 50]$ s, $f_\omega = 4000$ 1/min during $t \in [60, 70]$ s and $f_{p_{\mathrm{p,out}}} = 10$ bar during $t \in [80, 90]$ s. As Hybrid and SVN achieve the highest isolation performance, their qualitative diagnostic output over time for the test run is shown in

Table 6 Performance on the unknown fault class. The best value for each fault class is highlighted in bold.

System	Metric	$f_{\mathrm{BPV}} = 0.5$	$f_{\mathrm{BPV}} = 0.5$	$f_{\mathrm{BPV}} = 0.6$
Hybrid	$\mathrm{FAR}_{\mathrm{unk}}$	**0.00**	**0.00**	**0.01**
CBD	$\mathrm{FAR}_{\mathrm{unk}}$	**0.00**	**0.00**	**0.01**
SVN	$1 - \mathrm{IA}_{\mathrm{unk}}$	0.02	0.02	0.02

Figure 7. Within the first nominal segment $t \in [0, 20]$ s, both architectures correctly detect nominal operation. When f_η is introduced, the residuals of both architectures deviate from zero. SVN correctly diagnoses this fault after the warm-up phase of the persistence filter, while Hybrid fails to isolate it correctly due to the overlap between f_η and $f_{p_{\mathrm{t,in}}}$ in the residual space (see Section 5.1). In the subsequent nominal segment, the return to a nominal diagnosis is briefly delayed for both architectures as a result of the persistence filter. When the unknown fault f_{BPV} is introduced, the SVN residuals deviate from zero, and the unknown fault class is correctly diagnosed after the warm-up phase of the persistence filter. Hybrid does not respond to this fault class. As a result, no warm-up phase is required, and the correct diagnosis is maintained throughout. Again, both architectures correctly classify nominal operation in the following segment. When f_ω is introduced, both architectures correctly diagnose the fault after the warm-up phase of the persistence filter. Notably, the residual response of the hybrid architecture is stronger for f_ω than for f_η, indicating a higher fault sensitivity. The following nominal segment is again correctly classified by both architectures. In the final fault segment, $f_{p_{\mathrm{p,out}}}$ is correctly diagnosed by both architectures after the warm-up phase of the persistence filter.

Exploiting differences between sensor- and component-fault signatures, SVN correctly and unambiguously isolates all fault episodes, including structurally indistinguishable and unknown classes, where Hybrid and CBD reach their limits.

5.4 Generalizability

The scope of the proposed approach is generalization across LPREs rather than across arbitrary complex systems. Within this scope, the methodology's requirements are satisfied by construction. The governing equations of the reduced-order model follow from first-principles conservation laws and are therefore known. Validated transient simulation models are firmly embedded in the engine development process and thus available for FDI design [25]. The dominant fault mechanisms of LPREs are likewise well documented [4, 14], providing the basis for fault modeling in both the structural analysis and the generation of training data. Since the OC-SVM and SVN training data are also obtained from the validated simulation model, the data-driven components inherit the same availability as the structural model. The methodology can thus be adapted to other engine configurations by exchanging the underlying simulation model and training data, without modifying the diagnostic architecture.

Extending the approach from the FTP subsystem to the full engine follows the same workflow. The structural model is augmented with the governing equations of the remaining components, the detectability and isolability analysis is repeated for the extended fault set and sensor configuration and the data-driven components are retrained on correspondingly generated data. As structural isolability depends on the sensor configuration and fault set, the attainable diagnostic performance must be re-assessed for each target system.

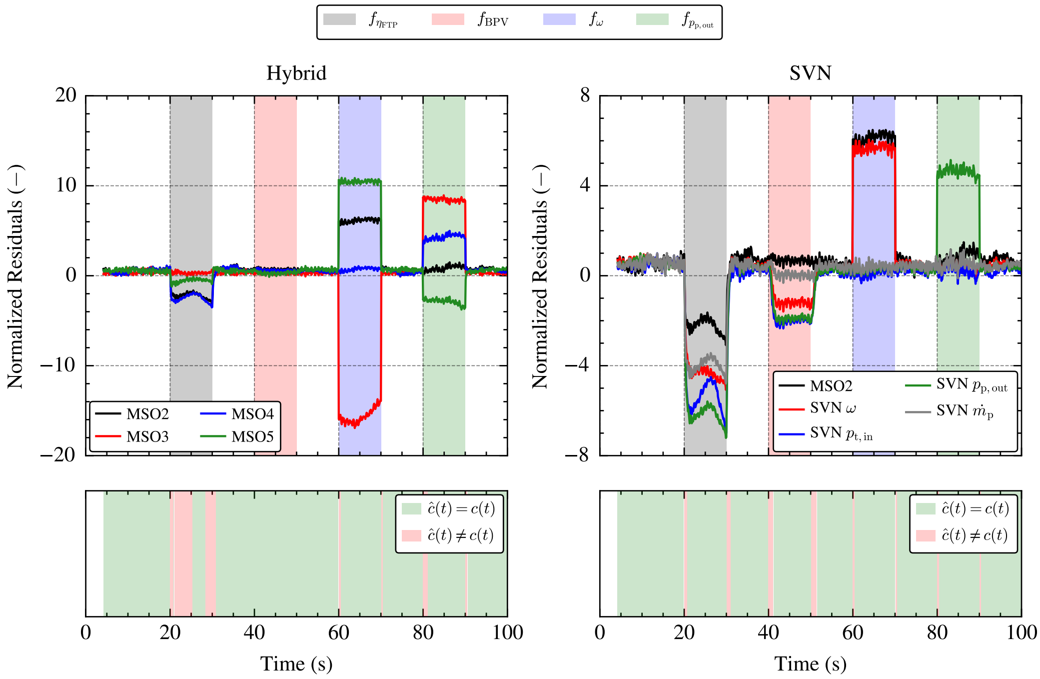

Figure 7 Trajectory-level performance of the hybrid architecture and SVN for multiple faults. Each residual is normalized using the fault-free first 20 s of the trajectory.

6 Conclusion and Outlook

In this work, we demonstrated that incorporating data-driven components into a model-based framework for fault detection and isolation improves the performance for the FTP subsystem of the LOX/LNG rocket engine. In the residual-generation stage, we selectively employ neural ODE residual generators for MSO sets for which direct model-based residual generation would require numerically ill-conditioned inversions. By parameterizing the system dynamics in grey-box form, additional MSO sets can be made accessible, thereby maximizing structural isolability. To address the structural indistinguishability between two of the five considered faults, we employ a data-driven isolation layer based on a set of one-class classifiers. Each classifier is trained on residual data from a single fault class, enabling the diagnosis system to account for distributional differences not captured by structural fault signatures alone. The hybrid architecture substantially improves isolation accuracy (IA) compared with the consistency-based baseline, increasing the IA from 0.46 to 0.76.

We further demonstrate that exploiting differences between sensor- and component-fault signatures yields an additional improvement in diagnostic performance, achieving an isolation accuracy of 0.92. For this purpose, we employed a grey-box neural ODE surrogate driven by control-command inputs, which serves as a sensor validation network (SVN). This fault decoupling enables local structural analysis for component-fault diagnosis, resulting in targeted diagnostic modules for selected component faults. The SVN architecture further improves the isolation accuracy to 0.92, while retaining the ability to detect unknown component faults.

Future work will extend our proposed approach to a distributed diagnosis system that covers all LUMEN subsystems. In addition, the approach will be demonstrated under closed-loop engine operation and extended to gradual fault drifts. The simulation-to-reality gap will be addressed by incorporating slow thermal transients into the neural ODE models and adapting the surrogates to experimental data. Finally, the performance of the resulting FDI system will be validated with experimental data.

References

1 J. Cha. Numerical Simulation of Chemical Propulsion Systems: Survey and Fundamental Mathematical Modeling Approach. *Aerospace*, 10(10), 2023. doi:10.3390/aerospace10100839.

2 J. Cha and S. Ko. General Fault Factor Method for Health Monitoring of a Liquid Rocket Engine. *Journal of Spacecraft and Rockets*, 62(6):2236–2248, 2025. doi:10.2514/1.A36337.

3 R. L. Christenson, M. A. Nelson, and J. P. Butas. Rocket Engine Health Management: Early Definition of Critical Flight Measurements. In *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, Huntsville, Alabama, 2003. doi:10.2514/6.2003-5251.

4 H. A. Cikanek. Characteristics of Space Shuttle Main Engine Failures. In *23rd Joint Propulsion Conference*, San Diego, California, June 1987. AIAA Paper 87-1939. doi:10.2514/6.1987-1939.

5 M. Davidson and J. Stephens. Advanced Health Management System for the Space Shuttle Main Engine. In *40th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, Fort Lauderdale, FL, July 2004. AIAA 2004-3912.

6 J. de Kleer. Hitting Set Algorithms for Model-Based Diagnosis. In *22nd International Workshop on Principles of Diagnosis (DX'11)*, Murnau, Germany, 2011.

7 J. de Kleer and B. C. Williams. Diagnosing Multiple Faults. *Artificial Intelligence*, 32(1):97–130, 1987. doi:10.1016/0004-3702(87)90063-4.

8 K. Dresia, E. Kurudzija, V. Bareiß, J. Dauer, J. Deeken, and G. Waxenegger-Wilfing. Health Monitoring and Intelligent Control Benchmark for Reusable Liquid-Propellant Rocket Engines. In *Space Propulsion Conference*, Bari, Italy, 2026. 3AF. doi:10.60711/SPC2026.20260715.37863079794095352.

9 K. Dresia, T. Traudt, M. Börner, D. Suslov, W. Armbruster, R. dos Santos Hahn, E. Kurudzija, C. Groll, M. A. Müller, S. Klein, J. Hämisch, J. Deeken, J. Hardi, and S. Schlechtriem. Hot-Fire Testing and System Analysis of the LUMEN Liquid Upper Stage Demonstrator Engine. In *3rd Int. Conf. on Flight Vehicles, Aerothermodynamics and Re-entry (FAR)*, Arcachon, France, 2025.

10 European Commission. ENLIGHTEN: European iNitiative for Low Cost, Innovative & Green High Thrust Engine, October 2022. doi:10.3030/101082326.

11 B. Férard, S. Le Gonidec, M. Galeotta, S. Oriol, and S. Dreyer. Anomaly Detection on Propulsive Systems by Global Approach Using Autoencoders. *IFAC-PapersOnLine*, 54(4):31–37, 2021. doi:10.1016/j.ifacol.2021.10.006.

12 E. Frisk, M. Krysander, and D. Jung. A Toolbox for Analysis and Design of Model Based Diagnosis Systems for Large Scale Models. *IFAC-PapersOnLine*, 50(1):3287–3293, 2017. 20th IFAC World Congress. doi:10.1016/j.ifacol.2017.08.504.

13 N. T. Ho. Failure Detection and Isolation for the Space Shuttle Main Engine. M.sc. thesis, Massachusetts Institute of Technology, 1998.

14 M. M. Hussein, D. Terakado, and O. de Weck. Comprehensive Study of the International Space Launch Industry: Programmatic Analysis and Technical Failures. *Acta Astronautica*, 234:676–689, 2025. doi:10.1016/j.actaastro.2025.03.015.

15 D. Jung. Data-Driven Open-Set Fault Classification of Residual Data Using Bayesian Filtering. *IEEE Transactions on Control Systems Technology*, 28(5):2045–2052, 2020. doi:10.1109/TCST.2020.2997648.

16 D. Jung. Automated Design of Grey-Box Recurrent Neural Networks For Fault Diagnosis using Structural Models and Causal Information. *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, 168:8–20, 23–24 June 2022. URL: https://proceedings.mlr.press/v168/jung22a.html.

17 D. Jung, K. Y. Ng, E. Frisk, and M. Krysander. Combining Model-Based Diagnosis and Data-Driven Anomaly Classifiers for Fault Isolation. *Control Engineering Practice*, 80:146–156, 2018. `doi:10.1016/j.conengprac.2018.08.013`.

18 D. Jung and T. Westny. Uncertainty-Aware Fault Diagnosis of Unknown Faults using Ensemble-Based NODE Residuals. *Mechanical Systems and Signal Processing*, 242:113599, 2026. `doi:10.1016/j.ymssp.2025.113599`.

19 C. Keliris, M. M. Polycarpou, and T. Parisini. A Distributed Fault Detection Filtering Approach for a Class of Interconnected Continuous-Time Nonlinear Systems. *IEEE Transactions on Automatic Control*, 58(8):2032–2047, 2013. `doi:10.1109/TAC.2013.2253231`.

20 H. Khorasgani and G. Biswas. Structural Fault Detection and Isolation in Hybrid Systems. *IEEE Transactions on Automation Science and Engineering*, 15(4):1585–1599, 2017. `doi:10.1109/TASE.2017.2749447`.

21 H. Khorasgani, A. Farahat, K. Ristovski, C. Gupta, and G. Biswas. A Framework for Unifying Model-Based and Data-Driven Fault Diagnosis. In *Proceedings of the Annual Conference of the PHM Society*, volume 10, September 2018. `doi:10.36001/phmconf.2018.v10i1.530`.

22 M. Krysander, J. Åslund, and M. Nyberg. An Efficient Algorithm for Finding Minimal Overconstrained Subsystems for Model-Based Diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 38(1):197–206, 2008. `doi:10.1109/TSMCA.2007.909555`.

23 E. Kurudzija, K. Dresia, V. Bareiß, G. Waxenegger-Wilfing, and J. Deeken. Transient System-Level Simulation of the LOX/LNG Expander-Bleed Rocket Engine LUMEN. SSRN. https://ssrn.com/abstract=6806858, 2026.

24 K. Lee, J. Cha, S. Ko, S.-Y. Park, and E. Jung. Fault Detection and Diagnosis Algorithms for an Open-Cycle Liquid Propellant Rocket Engine using the Kalman Filter and Fault Factor Methods. *Acta Astronautica*, 150:15–27, 2018. `doi:10.1016/j.actaastro.2018.03.001`.

25 M. A. Martin, H. H. Nguyen, W. D. Greene, and D. C. Seymout. Transient Mathematical Modeling for Liquid Rocket Engine Systems: Methods, Capabilities, and Experience. In *5th International Symposium on Liquid Space Propulsion*, Chattanooga, Tennessee, October 2003.

26 A. Mohammadi, M. Krysander, and D. Jung. Analysis of Grey-Box Neural Network-Based Residuals for Consistency-Based Fault Diagnosis. *IFAC-PapersOnLine*, 55(6):1–6, 2022. `doi:10.1016/j.ifacol.2022.07.097`.

27 A. Mohammadi, M. Krysander, and D. Jung. Consistency-Based Diagnosis Using Data-Driven Residuals and Limited Training Data. *Control Engineering Practice*, 159:106283, 2025. `doi:10.1016/j.conengprac.2025.106283`.

28 R. Murata, J. Marzat, H. Piet-Lahanier, S. Boujnah, and P. Belleoud. Residual Selection for Observer-Based Fault Detection and Isolation in a Multi-Engine Propulsion Cluster. In *Proceedings of the European Conference of the PHM Society*, volume 8, page 9, June 2024. `doi:10.36001/phme.2024.v8i1.4056`.

29 R. Murata, J. Marzat, H. Piet-Lahanier, S. Boujnah, and P. Belleoud. Fault Detection and Isolation Based on Structural Analysis: Application to a Multi-Engine Propulsion Cluster. *Sensors*, 25(4), 2025. `doi:10.3390/s25041054`.

30 J. W. Owen. Shuttle Propulsion System Major Events and the Final 22 Flights. In *6th JANNAF Liquid Propulsion Subcommittee Meeting, NASA Marshall Space Flight Center*, Huntsville, AL, USA, 2011.

31 S.-Y. Park and J. Ahn. Deep Neural Network Approach for Fault Detection and Diagnosis During Startup Transient of Liquid-Propellant Rocket Engine. *Acta Astronautica*, 177:714–730, 2020. `doi:10.1016/j.actaastro.2020.08.019`.

32 I. Pill, D. Jung, E. Kurudzija, A. Sztyber-Betley, M. Syfert, K. Dresia, G. Waxenegger-Wilfing, and J. de Kleer. The DX Competition 2025 and Its Benchmarks. *Open Access Series in Informatics (OASIcs)*, 136:14:1–14:19, 2025. 36th International Conference on Principles of Diagnosis and Resilient Systems (DX 2025). `doi:10.4230/OASIcs.DX.2025.14`.

33 S. Priotto and S. Le Gondiec. AI-Driven Health Monitoring Systems for Enhanced Propulsion Reliability. In *Space Propulsion Conference*, Bari, Italy, 2026. 3AF.

34 R. Reiter. A Theory of Diagnosis from First Principles. *Artificial Intelligence*, 32(1):57–95, 1987. doi:10.1016/0004-3702(87)90062-2.

35 B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001. doi:10.1162/089976601750264965.

36 P. Simontacchia, S. Sagnier, E. Edeline, R. Blasi, A. Espinosa, J.-N. Caruana, S. Boujnah, and P. Altenhöfer. PROMETHEUS: Precursor of New Low-Cost Rocket Engine Family. In *Proceedings of the International Astronautical Congress, IAC*, 2022.

37 K. Tidriri, N. Chatti, S. Verron, and T. Tiplica. Bridging Data-Driven and Model-Based Approaches for Process Fault Diagnosis and Health Monitoring: A Review of Researches and Future Challenges. *Annual Reviews in Control*, 42:63–81, 2016. doi:10.1016/j.arcontrol.2016.09.008.

38 T. Traudt, W. Armbruster, C.r Groll, R. H. Dos Santos Hahn, K. Dresia, M. Börner, S. Klein, D. Suslov, E. Kurudzija, J. Haemisch, M. A. Müller, J. C. Deeken, J. Hardi, and S. Schlechtriem. LUMEN, the Test Bed for Rocket Engine Components: Results of the Acceptance Tests and Overview on the Engine Test Preparation. In *9th Edition of the Space Propulsion Conference*, Glasgow, Scotland, 2024.

39 C. Wang, Y. Zhang, Z. Zhao, X. Chen, and J. Hu. Dynamic Model-Assisted Transferable Network for Liquid Rocket Engine Fault Diagnosis Using Limited Fault Samples. *Reliability Engineering & System Safety*, 243:109837, 2024. doi:10.1016/j.ress.2023.109837.

40 T. Wang, L. Ding, and H. Yu. Research and Development of Fault Diagnosis Methods for Liquid Rocket Engines. *Aerospace*, 9(9):481, 2022.

41 J. Wu. Liquid-Propellant Rocket Engines Health-Monitoring — a Survey. *Acta Astronautica*, 56(3):347–356, 2005. doi:10.1016/j.actaastro.2004.05.070.

42 X. Zhu, Y. Cheng, J. Wu, R. Hu, and X. Cui. Steady-State Process Fault Detection for Liquid Rocket Engines Based on Convolutional Auto-Encoder and One-Class Support Vector Machine. *IEEE Access*, 8:3144–3158, 2020. doi:10.1109/ACCESS.2019.2961742.