electronics

MDPI

Article

Physics-Informed Co-Optimization of Fuel-CellFlying Vehicle Propulsion and Control Systems with Onboard Catalysis

Yifei Bao [1], Chaoyi Chen [2], Hao Zhang [3,*] and Nuo Lei [2,*]

[1] School of Engineering, Newcastle University, Newcastle NE1 7RU, UK
[2] School of Vehicle and Mobility, Tsinghua University, Beijing 100084, China
[3] Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
* Correspondence: haoz4@andrew.cmu.edu (H.Z.); lein21@mails.tsinghua.edu.cn (N.L.)

Abstract

Fuel-cell flying vehicles suffer from limited endurance, while ammonia, decomposed onboard to supply hydrogen, offers a carbon-free, high-density solution to extend flight missions. However, the system's performance is governed by a multi-scale coupling between propulsion and control systems. To this end, this paper introduces a novel optimization paradigm, termed physics-informed gradient-enhanced multi-objective optimization (PI-GEMO), to simultaneously optimize the ammonia decomposition unit (ADU) catalyst composition, powertrain sizing, and flight control parameters. The PI-GEMO framework leverages a physics-informed neural network (PINN) as a differentiable surrogate model, which is trained not only on sparse simulation data but also on the governing differential equations of the system. This enables the use of analytical gradient information extracted from the trained PINN via automatic differentiation to intelligently guide the evolutionary search process. A comprehensive case study on a flying vehicle demonstrates that the PI-GEMO framework not only discovers a superior set of Pareto-optimal solutions compared to traditional methods but also critically ensures the physical plausibility of the results.

Keywords: physics-informed optimization; system coupling; fuel-cell flying vehicle; gradient-enhanced search; Pareto-optimal solutions

check for updates

Academic Editor: Andrea Bonci

Received: 14 September 2025
Revised: 14 October 2025
Accepted: 21 October 2025
Published: 23 October 2025

Citation: Bao, Y.; Chen, C.; Zhang, H.; Lei, N. Physics-Informed Co-Optimization of Fuel-CellFlying Vehicle Propulsion and Control Systems with Onboard Catalysis. *Electronics* **2025**, *14*, 4150. https://doi.org/10.3390/electronics14214150

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The advent of flying vehicles signals a potential paradigm shift in urban and regional transportation, promising to alleviate terrestrial traffic congestion, accelerate logistics, and enable a host of novel aerial services [1]. Despite significant progress in aerodynamics, materials science, and autonomous control, the widespread deployment of flying vehicles is fundamentally constrained by a persistent and critical challenge: the design of the high-energy-density propulsion system. The current state-of-the-art energy source for electric vertical take-off and landing (eVTOL) and flying vehicles, the lithium-ion battery, possesses a gravimetric energy density [2]. While sufficient for short-duration flights, this value falls dramatically short of the estimated 1000–1200 Wh/kg required for missions of commercially viable range and payload capacity [3]. The energy-intensive takeoff and hover phases of flight rapidly deplete stored energy, relegating most current designs to intracity operations and limiting their economic feasibility. Furthermore, issues of long charging times and battery degradation over cycle life present additional significant barriers to the scaling pf flying vehicle operations [4].

In response to the limitations of battery-only propulsion, the research community has begun exploring alternative energy vectors to extend the operational envelope of electric

aviation. Gaseous hydrogen, when used in conjunction with a proton-exchange membrane fuel cell (PEMFC), offers an exceptional gravimetric energy density (33.3 kWh/kg, LHV), far exceeding that of batteries [5]. PEMFC is superior in terms of of cleanliness, with low noise and high energy density, and is considered the preferred energy source for transportation vehicles [6]. However, the limited lifespan of PEMFC still hinders its commercialization. Specifically, fuel cells experience a gradual decline in performance during long-term operation. Performance degradation not only affects the efficiency and output power of fuel cells but also shortens their service life and increases maintenance and replacement costs [7]. Moreover, their practical application in mass-sensitive aerial vehicles is severely hampered by their extremely low volumetric density. Storing hydrogen onboard requires either heavy, high-pressure composite tanks (e.g., 700 bar) or complex and energy-intensive cryogenic liquid storage systems [8]. Amidst these competing technologies, ammonia (NH3) has emerged as a highly compelling carbon-free hydrogen carrier and energy vector for propulsion applications [9]. As a molecule, ammonia possesses a high gravimetric (17.6 wt%) and exceptional volumetric hydrogen density, enabling it to be stored as a liquid at moderate pressures [10]. Moreover, ammonia benefits from a well-established global infrastructure for its production, transportation, and storage, which could be adapted to support a future energy economy [11]. However, the transition from a conceptual powertrain to a functional, optimized system for an eVTOL introduces a design challenge of formidable complexity.

An ammonia-powered aerial vehicle is not merely a collection of components but a deeply integrated, multi-physics system characterized by strong, bidirectional couplings across multiple domains and scales [12]. A critical chemical–thermal coupling exists between the ammonia decomposition unit (ADU) and the PEMFC [13]. The catalytic decomposition of ammonia is an endothermic process, requiring a continuous heat supply to maintain the catalyst at its optimal operating temperature (typically >400 °C) [14]. The performance of the catalyst is a non-linear function of temperature, which, in turn, depends on the power output and, thus, the efficiency of the fuel cell [15]. Furthermore, a tight power-dynamic coupling links the powertrain's output to the vehicle's flight mechanics. Moreover, the control strategy for the vehicle's flight path directly dictates this power profile, meaning the controller's design parameters are intrinsically linked to the energy consumption of the powertrain [16]. In addition, for an aerial application, a pervasive mass-performance coupling dominates the entire design. Unlike ground vehicles, where mass is a secondary consideration [17], components such as the ADU reactor, the PEMFC stack, the motors, and the battery pack contribute directly to the gross takeoff mass [18]. Addressing such a tightly coupled system necessitates a holistic, co-design optimization approach [19]. Traditional sequential design methodologies, where the powertrain is first optimized for a static power profile and a flight controller is subsequently designed for the resulting hardware, are inadequate, as they fail to capture the critical feedback loops between the subsystems [20]. A feasible co-design approach should simultaneously optimize variables from all domains: microscopic catalyst composition in the ADU, macroscopic sizing of the PEMFC and battery, and control-law parameters for the flight management system. Direct optimization using high-fidelity models is therefore computationally intractable, often requiring years of compute time. This computational barrier necessitates the use of surrogate models—computationally cheap mathematical approximations of the high-fidelity model [21].

One prevailing paradigm in engineering design optimization is the use of data-driven surrogate models. Techniques such as Gaussian processes, polynomial chaos expansion, and standard artificial neural networks have been widely employed [22]. These methods function by learning a statistical mapping from input design parameters to output per-

formance metrics based on a limited set of high-fidelity data points [23]. While effective in many contexts, their application to complex, safety-critical systems like flying vehicles reveals a fundamental flaw [24]. As black-box models, they possess no intrinsic knowledge of the underlying physics governing the system [25]. To overcome this critical limitation, a paradigm shift in surrogate modeling is required, moving from purely data-driven methods to those that are physics-informed. Physics-informed neural networks (PINNs) have recently emerged as a powerful tool for this purpose [26]. A PINN is a neural network that is trained to minimize a composite loss function. This function includes not only the conventional data-misfit term but also a physics-based residual term that penalizes deviations from the system's governing ordinary or partial differential equations [27]. This physics loss acts as a powerful, continuous form of regularization, constraining the network's output to lie on the manifold of physically admissible solutions [28]. While the adoption of a PINN as a surrogate represents a significant advancement, this paper argues that its full potential is unlocked only when its unique properties are leveraged to innovate the optimization algorithm itself [29]. A trained PINN is not just a fast and accurate calculator but a fully differentiable, analytical representation of the complex input–output relationship of the system [30]. This differentiability, enabled by the automatic differentiation capabilities inherent in modern deep learning frameworks, allows for the instantaneous and precise computation of the gradient of any performance metric with respect to any design parameter [31]. However, it is crucial to acknowledge that the 'physical plausibility' enforced by a PINN is inherently limited by the completeness of the governing equations. Unmodeled physical phenomena, such as long-term catalyst sintering deactivation or complex fuel-cell water management dynamics, represent a challenge that requires explicit declaration of the model's boundary conditions and scope [32].

This work aims to fill this methodological gap by proposing a novel optimization framework that deeply integrates the differentiable nature of PINNs into the search process. This paper introduces physics-informed gradient-enhanced multi-objective (PI-GEMO) optimization, a new class of hybrid intelligent optimization designed explicitly for complex, multi-physics engineering systems. PI-GEMO synergistically fuses two distinct search strategies within a single evolutionary loop. It retains the powerful, gradient-free global exploration operators of traditional genetic algorithms to ensure population diversity and prevent premature convergence. Critically, it introduces a novel gradient-guided local search operator that uses the analytical gradients extracted from the PINN to perform efficient, multi-objective gradient descent on promising candidate solutions, rapidly propelling them towards the true Pareto-optimal front. The main contributions of this paper include the following: (1) A high-fidelity, multi-physics dynamic model of an ammonia-powered fuel cell flying vehicle is developed, capturing the tight couplings between the chemical, thermal, electrical, and aerodynamic domains. (2) PI-GEMO leverages a differentiable PINN surrogate to create a hybrid algorithm that synergizes gradient-free global exploration with gradient-based local search for enhanced efficiency and precision. (3) Comprehensive tests demonstrate the application of the proposed PI-GEMO, with comparisons with conventional approaches, and provide critical insights into the co-design trade-offs of propulsion and control systems for flying vehicles.

The remainder of this paper is organized as follows: Section 2 describes the flying vehicle powertrain model and control system. Section 3 formulates the proposed physics-informed co-optimization methodology. Section 4 discusses the results of coupled optimization. Conclusions are summarized in Section 5.

2. Flying Vehicle System Modeling

This section details the high-fidelity, multi-physics dynamic model of an ammonia-powered flying vehicle. The hybrid electric flying vehicle investigated in this study adopts a hexacopter configuration, driven by six independent electric motors. The propulsion system is supported by a dual-fuel-cell architecture, in which two PEMFC stacks operate in parallel to provide a stable and continuous power supply throughout the flight, which can be seen in Figure 1. The high-fidelity model serves a dual purpose: it is used to generate the sparse dataset for training of the surrogate model; in addition, its governing equations are embedded directly into the PINN's loss function to ensure the predictions adhere to physical laws. The model is decomposed into interconnected sub-models for the vehicle power demand, the ammonia–hydrogen propulsion system, and the power management system.

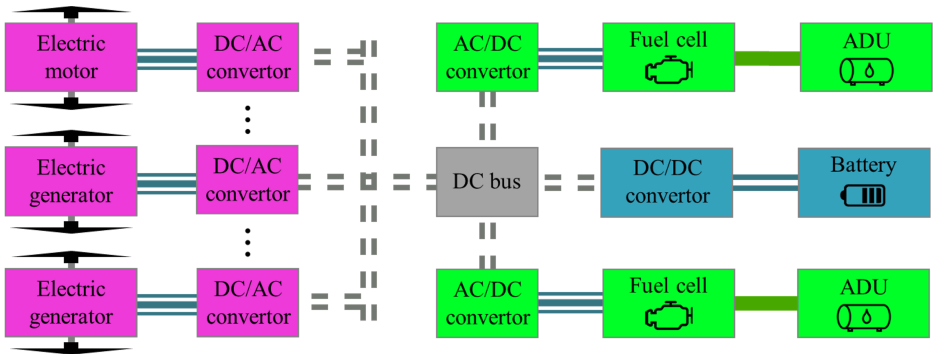


Figure 1. Drivetrain and powertrain configuration of the flying vehicle.

2.1. Vehicle Power Demand Model

To accurately quantify the instantaneous power required for flight, a quasi-steady-state aerodynamic model is employed, which captures both vertical and horizontal components of rotor power. Figure 2 depicts a schematic of the flying vehicle. The total rotor power demand (P_{demand}) is the sum of its vertical (P_V) and horizontal (P_H) components, which depend on the effective vertical (v_v) and horizontal (v_h) flight velocities:

$$\begin{cases} P_{demand}(t) = P_V(t) + P_H(t) \\ v_v = v_{air} \sin\phi \\ v_h = v_{air} \cos\phi \end{cases} \tag{1}$$

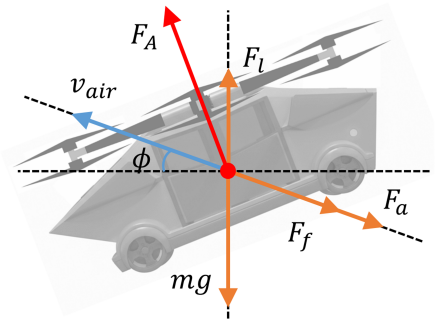


Figure 2. Quasi-steady-state dynamics model of the flying vehicle, showing forces and velocity components. The colored arrows refer to the forces and the dashed black lines refer to the axles.

The vertical component (P_V) reflects the balance between thrust and gravity during climbing or descent. Following rotorcraft theory [33], it can be expressed as follows:

$$P_V = \begin{cases} \frac{mg}{2} v_v + \frac{mg}{2} \sqrt{v_v^2 + \frac{2mg}{\rho \pi R_0^2}}, & \text{climbing} \\ \frac{mg}{2} v_v - \frac{mg}{2} \sqrt{v_v^2 - \frac{2mg}{\rho \pi R_0^2}}, & \text{descending} \end{cases} \tag{2}$$

where ρ is air density and R_0 is the rotor radius.

The horizontal component (P_H) represents the power required to overcome parasitic drag (P_P) and lift-induced effects (P_I). These components are calculated based on aerodynamic principles, where the induced inflow ratio (ψ) is found by implicitly solving a momentum theory equation [33]:

$$\begin{cases} P_H = P_P + P_I \\ P_P = \frac{1}{2} \rho C_{D,p} A_e v_h^3 + \frac{\pi}{4} N_b c_b \rho C_{d,b} \omega^3 R_0^4 \left(1 + 3 \left(\frac{v_h}{\omega R_0} \right)^2 \right) \\ P_I = mg \omega R_0 \psi \\ \psi \text{ is solved from: } 2\rho \pi \omega^2 R_0^4 \psi \left(\frac{v_h}{\omega R_0} \right)^2 + \psi^2 - mg = 0 \end{cases} \tag{3}$$

where $C_{D,p}$ is the parasitic drag coefficient, N_b is the number of blades, c_b is the blade chord length, $C_{d,b}$ is the blade-profile drag coefficient, and ω is rotor angular velocity. This quasi-steady-state approach is adopted to ensure the optimization problem remains computationally tractable, though we acknowledge it does not capture high-frequency transient effects such as gust disturbances.

2.2. Ammonia–Hydrogen Propulsion System

The propulsion system generates the required electrical power by converting chemical energy from ammonia. It consists of a catalytic ADU and a PEMFC stack, which are benchmarked and calibrated against experimental data from authoritative literature and manufacturer datasheets. It is critical to note that the high-temperature ADU and the low-temperature PEMFC are thermally decoupled systems. The ADU's catalyst bed is maintained at its high operating temperature (>400 °C) via an integrated electrical heater, while the PEMFC stack operates within its optimal lower temperature range (60–80 °C) using a separate liquid cooling loop. The mass of the ADU reactor, including the necessary high-temperature insulation, is accounted for in the overall system mass model.

2.2.1. Ammonia Decomposition Unit

The design and performance model of the ammonia decomposition unit is central to the entire propulsion system, whose structure can be found in Figure 3. The core of the system is a multi-catalyst reactor, where the key design variables for optimization are the composition, cost, and selection of catalysts. The performance maps and kinetic parameters used in this model were benchmarked against a comprehensive set of experimental data from authoritative literature to ensure physical fidelity. The unit costs for precious metals are based on the average market prices over a recent three-year period (2022–2024), referenced from major commodity market indices like the London Metal Exchange (LME), to provide a stable economic baseline. We acknowledge that this model simplifies costs by not including detailed carrier processing expenses or end-of-life degradation costs; these factors are noted as important considerations for future, more detailed techno-economic analyses.

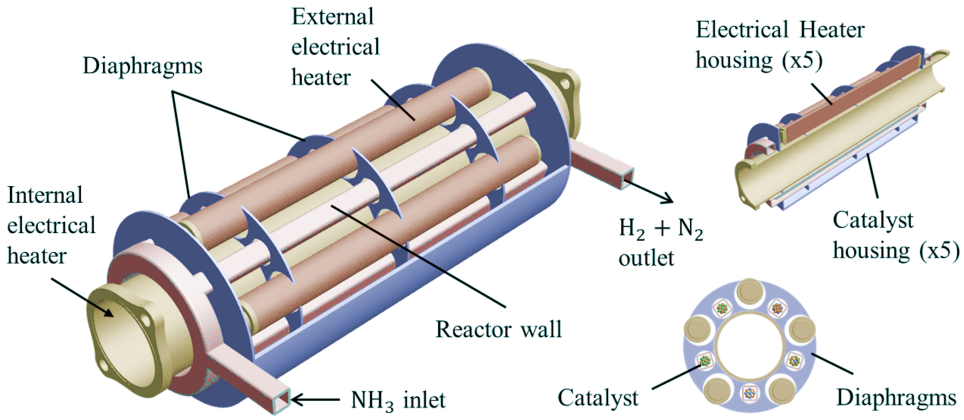


Figure 3. Schematic diagram of the physical structure of the ADU.

Catalyst Selection and Performance Characterization

The ADU's efficacy is fundamentally dictated by the judicious selection of its catalysts. The detailed properties and costs, summarized in Tables 1 and 2 respectively, are collated from a comprehensive review of recent experimental literature and manufacturer specifications. While the cited data provide a strong basis for our optimization, it is acknowledged that catalyst performance and precious metal costs can vary; the robustness of our final design choices to such uncertainties was confirmed via the sensitivity analysis presented in Section 4.3. In the ammonia decomposition reaction, the catalyst's role is to lower the activation energy required to break the strong N–H bonds. The detailed properties are summarized in Table 1 for specifications and Table 2 for costs [4]. These catalysts, ranging from precious metal-based to non-precious metal-based, exhibit diverse performance characteristics in terms of activity, stability, and cost [34].

Table 1. Catalyst specifications with values of main parameters [4,35].

No.	Catalyst (A-E/F/G-X/Y)	Temp (°C)	GHSV (mL/g/h)	x (%)	α (*w*/*w*%)
A	α	400–500	15,000	84–100	5
B	α	400–500	36,000	40–84	7.8
C	α	400–500	30,000	10–65	100
D	α	400–500	7200	17–100	100
E	α	400–500	30,000	18–72.4	60
F	α	400–500	13,800 (Ar:NH3 = 1.3:1)	14–95	10
G	α	400–500	13,800 (Ar:NH3 = 1.3:1)	77–100	2
H	α	400–500	6000	9–73.8	34.7
I	α	400–500	18,000	32–65	4.8
J	α	400–500	15,000	49.7–99.2	2.74

The central performance metric is the ammonia conversion efficiency (η_{conv}), which is a complex function of operating temperature (T_{ADU}), the specific catalyst mixture ($\boldsymbol{\rho}_{cat}$), and the gas hourly space velocity (GHSV) [36], which greatly simplifies the model, although it may understate rapid transient dynamics, which can be ignored in this work for the systematic design of the powertrain. GHSV is defined as the ratio of the volumetric ammonia flow rate ($\dot{V}_{NH_3}$) to the catalyst volume (V_{cat}). These relationships are captured by the following equations:

$$\begin{cases} \text{GHSV} = \frac{\dot{V}_{NH_3}}{V_{cat}} \\ \eta_{conv} = f_{map,cat}(\boldsymbol{\rho}_{cat}, T_{ADU}, \text{GHSV}) \end{cases} \quad (4)$$

where the conversion efficiency is determined based on multi-dimensional lookup tables derived from empirical data and $\boldsymbol{\rho}_{cat}$ is a key set of decision variables in the optimization problem.

Catalyst Cost Model

The cost of the catalyst bed is a critical objective function in the optimization. The total cost is calculated based on the mass fraction ($\rho_{cat,i}$) and unit cost ($P_{cat,i}$) of each constituent component, including promoters, active materials, and supports:

$$\begin{cases} C_{cat,total} = m_{cat,total} \sum_{i=1}^{N_{cat}} \rho_{cat,i} \cdot P_{cat,i} \\ P_{cat,i} = P_{pro} W_{pro} + P_{act} W_{act} + P_{sup} W_{sup} \end{cases} \tag{5}$$

where $m_{cat,total}$ is the total catalyst mass and, in the second equation, P and W represent the unit cost and mass fraction of the promoter, active component, and support material.

Table 2. Costs of catalyst with components [4,37].

No.	Catalysts (A-E/F/G-X/Y)	α (w/w%)	β (CNY/g)	γ (w/w%)	Cost (CNY/g)	Cost (CNY/g)
A	α	1.38 (β = 3.62)	9.07	95	1134.83	66.63
B	α	7.8	55.54	92.2	33.92	53.85
C	α	100	-	0	13.36	13.36
D	α	100	-	0	5.44	5.44
E	α	60	45.13	40	2.39	19.49
F	α	10	2.39	90	45.13	6.66
G	α	2	2.39	98	1134.83	25.04
H	α	34.7	10.55	65.3	0.57	7.09
I	α	4.8	5.78	95.2	1134.83	59.97
J	α	2.8	2.39	97.2	180.6	7.38

Quasi-Steady-State Thermal Model

A quasi-steady-state thermal model is adopted for the ADU, assuming the temperature control system maintains a constant setpoint (T_{set}). Under this assumption, the instantaneous electrical power required by the heater ($P_{heater}(t)$) must balance the heat consumed by the endothermic reaction ($\dot{Q}_{react}(t)$) and the heat lost to the environment ($\dot{Q}_{loss}(t)$). This power balance represents a direct electrical load on the battery and is governed by the following:

$$\begin{cases} P_{heater}(t) = \dot{Q}_{react}(t) + \dot{Q}_{loss}(t) \\ \dot{Q}_{react}(t) = \dot{m}_{NH_3}(t) \cdot \eta_{conv}(t) \cdot \Delta H_{dec} \\ \dot{Q}_{loss}(t) = h_{conv}(v_{air}) A_{surf}(T_{set} - T_{amb}) \end{cases} \tag{6}$$

This quasi-steady-state model provides a robust and physically grounded mapping from the required hydrogen production rate to the necessary electrical heating power, forming a self-consistent and differentiable component model suitable for integration into the PINN framework.

2.2.2. Proton-Exchange Membrane Fuel-Cell System

To provide a detailed physical basis for the PINN, a high-fidelity, semi-empirical electrochemical model is employed instead of a simplified efficiency map. The model calculates the stack voltage by subtracting the primary irreversible voltage losses from the ideal thermodynamic potential (Nernst voltage). The output voltage of the fuel-cell stack (V_{stack}), which consists of N_{cell} individual cells connected in series, is given by the following

set of equations, which includes the Nernst voltage (E_{nernst}) and the voltage losses due to activation (V_{act}), ohmic resistance (V_{ohm}), and reactant concentration (V_{conc}):

$$\begin{cases} V_{\text{stack}}(t) = N_{\text{cell}}(E_{\text{nernst}}(t) - V_{\text{act}}(t) - V_{\text{ohm}}(t) - V_{\text{conc}}(t)) \\ E_{\text{nernst}} = 1.229 - 0.85 \times 10^{-3}(T_{\text{FC}} - 298.15) + 4.308 \times 10^{-5} T_{\text{FC}} \left[\ln(p_{\text{H}_2}) + \frac{1}{2}\ln(p_{\text{O}_2})\right] \\ V_{\text{act}} = \xi_1 + \xi_2 T_{\text{FC}} + \xi_3 T_{\text{FC}} \ln(C_{\text{O}_2}) + \xi_4 T_{\text{FC}} \ln(i) \\ V_{\text{conc}} = -B \ln\left(1 - \frac{i}{i_L}\right) \end{cases} \tag{7}$$

(1) Ohmic polarization (V_{ohm}), is the voltage drop due to resistance to the flow of protons through the polymer electrolyte membrane and electrons through the electrodes and current collectors. It is governed by Ohm's law. The membrane's proton conductivity (σ_m) is highly dependent on its water content (λ_m) and temperature:

$$\begin{cases} V_{\text{ohm}} = i \cdot R_{\text{ohm}} = i \cdot \frac{t_m}{\sigma_m} \\ \sigma_m = (0.005139\lambda_m - 0.00326)\exp\left[1268\left(\frac{1}{303} - \frac{1}{T_{\text{FC}}}\right)\right] \end{cases} \tag{8}$$

(2) The fuel-cell temperature (T_{FC}) is a critical state variable that influences all polarization losses. Its dynamic behavior is governed by the first law of thermodynamics, balancing heat generation and heat dissipation. The heat generation rate ($\dot{Q}_{\text{gen}}$) is the total chemical energy released minus the electrical work done, and heat dissipation ($\dot{Q}_{\text{diss}}$) primarily occurs through the coolant loop:

$$\begin{cases} m_{\text{FC}} C_{p,\text{FC}} \frac{dT_{\text{FC}}}{dt} = \dot{Q}_{\text{gen}}(t) - \dot{Q}_{\text{diss}}(t) \\ \dot{Q}_{\text{gen}}(t) = I_{\text{stack}}(t)\left(N_{\text{cell}} \frac{-\Delta H}{2F} - V_{\text{stack}}(t)\right) \\ \dot{Q}_{\text{diss}}(t) = \dot{Q}_{\text{FC,waste}}(t) = \dot{m}_{\text{coolant}} C_{p,\text{coolant}}(T_{\text{coolant,out}} - T_{\text{coolant,in}}) \end{cases} \tag{9}$$

This set of coupled, non-linear equations forms the high-fidelity physical model of the PEMFC, providing a detailed basis for the PINN's physics loss function.

2.2.3. Battery Model and Energy Management System

The battery pack is modeled using a zero-order Rint equivalent circuit model, a necessary trade-off for maintaining computational tractability at the mission-level optimization stage. The battery current (I_{bat}) is calculated according the battery power ($P_{bat}(t)$), and the state of charge (SOC) is then updated using the ampere-hour integration method. These two dynamics are described by the following coupled equations:

$$\begin{cases} I_{bat}(t) = \frac{V_{oc}(\text{SOC}) - \sqrt{V_{oc}(\text{SOC})^2 - 4R_{int}P_{bat}(t)}}{2R_{int}} \\ \frac{d(\text{SOC})}{dt} = -\frac{I_{bat}(t)}{C_{bat,rated}} \end{cases} \tag{10}$$

where the open-circuit voltage (V_{oc}) and internal resistance (R_{int}) are functions of the SOC and are determined from lookup tables, as shown in Figure 4. The internal resistance also depends on whether the battery is charging or discharging. $C_{bat,rated}$ is the rated capacity of the battery in ampere-hours.

The HEVTOL in this paper is equipped with the adaptive equivalent consumption minimization strategy (A-ECMS) for its energy management. A-ECMS is an optimization-based control method that minimizes fuel consumption by converting electrical energy into an equivalent fuel cost using an equivalence factor (EF), denoted by λ. The goal is to find the optimal battery power output (P_{bat}) at each moment to meet the total power demand (P_{PL}).

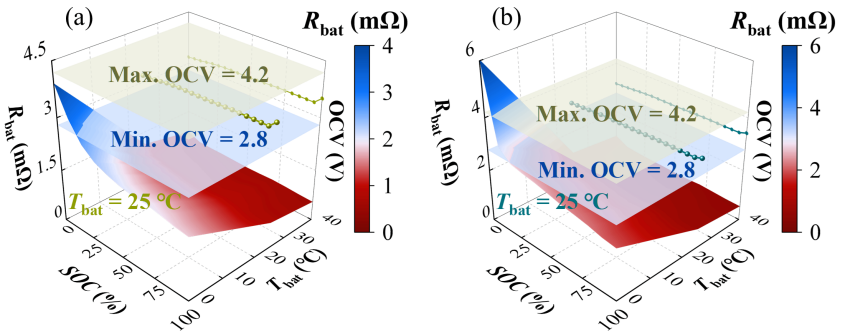


Figure 4. Battery characteristics as a function of SOC and temperature during during charging (**a**) and discharging (**b**).

The core of the strategy is to minimize an instantaneous cost function, the Hamiltonian, which represents the total equivalent fuel rate. This rate is the sum of the engine's actual fuel consumption ($\dot{m}_{\text{fuel}}$) and the equivalent fuel consumption of the battery power. To ensure the battery's SOC tracks a reference value (SOC_{ref}), the equivalence factor is made adaptive. A proportional–integral (PI) controller dynamically updates the EF at each time step (λ_t) based on the SOC deviation.

The entire A-ECMS optimization problem, which aims to minimize the cumulative equivalent fuel consumption over a mission, can be summarized as follows [38]:

$$\begin{cases} \underset{P_{\text{bat}}}{\text{Minimize}} & J = \int_{t_0}^{t_f} \left(\dot{m}_{\text{fuel}}(P_{\text{ICE}}) + \frac{\lambda_t}{LHV} P_{\text{bat}} \right) dt \\ \text{where} & \lambda_t = \lambda_0 + K_p(SOC - SOC_{\text{ref}}) + K_i \int (SOC - SOC_{\text{ref}})\, d\tau \\ \text{Subject to} & P_{\text{bat}} + P_{\text{ICE}} = P_{\text{PL}} \quad \text{and other physical limits (e.g., } SOC_{\min}, P_{\text{bat}}^{\max}) \end{cases} \tag{11}$$

In this unified framework, the objective is to minimize the total equivalent fuel burn (J). The adaptive mechanism is the PI controller that updates λ_t. The primary constraint is that the power from the battery and the engine (P_{ICE}) must meet the total demand (P_{PL}). The key control parameters to be optimized are the initial equivalence factor (λ_0) and the PI gains (K_p and K_i).

2.3. Power-Source State-of-Health Model

The durability of the power source in heavy-duty fuel-cell vehicles is crucial, as power fluctuations during operation significantly affect the lifespan of both the fuel cell and the power battery. To quantify their degradation, this section establishes semi-empirical state-of-health (SOH) models for both components.

2.3.1. Fuel-Cell SOH Model

The performance degradation of a fuel cell primarily stems from material aging and improper operating conditions. Key accelerated aging conditions include start–stops, load fluctuations, low-power operation, and high-power operation. The SOH degradation amount (ΔSOH_{FC}) can be evaluated by the following equation [39]:

$$\Delta SOH_{FC} = \psi_{cor}(\xi_{ss} n_{ss} + \xi_{lc} n_{lc} + \xi_{ll} t_{ll} + \xi_{hl} t_{hl}) \tag{12}$$

where ψ_{cor} is the correction coefficient; ξ_{ss}, ξ_{lc}, ξ_{ll}, and ξ_{hl} are the degradation rates corresponding to start–stop, load change, low-power, and high-power conditions, respectively; n_{ss} and n_{lc} are the number of start–stops and load changes, respectively; and t_{ll} and t_{hl} are the operating times under low and high power, respectively. The values for each parameter are shown in Table 3.

Table 3. Parameters for the fuel-cell degradation model [39].

Parameter and Symbol	Value and Unit
Correction coefficient (ψ_{cor})	1.47
Start–stop decay rate (ξ_{ss})	0.00196%/cycle
Load-change decay rate (ξ_{lc})	0.0000593%/cycle
Low-load decay rate (ξ_{ll})	0.00126%/h
High-load decay rate (ξ_{hl})	0.00147%/h

2.3.2. Battery SOH Model

Frequent charging/discharging and extreme temperatures accelerate battery aging. A semi-empirical model for its capacity loss (Q_{loss}^{bat}) is expressed as follows, where the α and β coefficients depend on the current state of charge (SOC):

$$Q_{loss}^{bat} = (\alpha \cdot \mathrm{SOC} + \beta) e^{\frac{-31700-163.3I_c}{RT_{bat}}} \mathrm{Ah}^{0.57} \quad \text{where} \begin{cases} \alpha = 1287.6,\ \beta = 6356.3 & \text{if SOC} \leq 0.45 \\ \alpha = 1385.5,\ \beta = 4193.2 & \text{if SOC} > 0.45 \end{cases} \tag{13}$$

where I_c is the charge/discharge rate, T_{bat} is the battery temperature, Ah is the ampere-hour throughput, and R is the molar gas constant.

To account for the difference between actual and nominal operating conditions, a severity factor (σ_{bat}) is introduced, and the effective ampere-hour throughput (Ah_{eff}) is defined as follows:

$$\sigma_{bat} = \frac{\Gamma_{nom}}{\Gamma_{real}} = \frac{\int_0^{L_{bat}} |I_{nom}| dt}{\int_0^{L_{bat}} |I_t| dt}, \quad Ah_{eff} = \int_{t_0}^{t_f} \sigma_{bat} |I_t| dt \tag{14}$$

The end of the battery's life is defined as occurring when its capacity degrades to 80%. At this point, the total life-cycle ampere-hour throughput (Γ_{20}) can be calculated using Equation (6) under nominal parameters ($I_{nom}^c = 2.5, \mathrm{SOC}_{nom} = 35\%$, and $T_{nom}^{bat} = 25\ ^\circ\mathrm{C}$). Finally, the SOH degradation of the battery can be expressed as the ratio of the effective ampere-hour throughput to the total throughput:

$$\Delta SOH_{bat} = \frac{Ah_{eff}}{\Gamma_{20}} \tag{15}$$

3. The PI-GEMO Optimization Framework

To solve the complex co-design problem defined by the tightly coupled, multi-physics model presented in Section 2, a novel optimization framework is proposed, as shown in Figure 5. The inherent conflict between the maximization of flight endurance and the minimization of system mass, coupled with a high-dimensional design space and computationally expensive simulations, necessitates a specialized approach. This section first formulates the co-design task as a multi-objective optimization problem (MOP). Subsequently, it details the architecture of our proposed PI-GEMO, a framework designed to synergize the global search capabilities of evolutionary algorithms with the efficiency of gradient-based methods, enabled by a bespoke physics-informed surrogate model.

3.1. Problem Formulation

This study focuses on the co-optimization of the ammonia–electric propulsion system's component sizing and the ADU's catalyst design. Let the optimization parameters collectively form a decision vector ($\mathbf{x} \in \mathbb{R}^N$). We aim to simultaneously minimize three conflicting objectives under the real-world driving scenarios shown in Figure 6. The objectives include the hydrogen consumption (J_{fuel}), the power-source degradation (J_{SOH}),

and the total powertrain cost J_{cost}. Combining the objectives, variables, and constraints, the co-design task is formally stated as the following MOP:

$$\begin{aligned} \min_{\mathbf{x}\in\Omega} \quad & \mathbf{F}(\mathbf{x}) = [J_{fuel}(\mathbf{x}), J_{SOH}(\mathbf{x}), J_{cost}(\mathbf{x})]^T \\ \text{s.t.} \quad & g_i(\mathbf{x}) \leq 0, \quad i = 1, \ldots, C \end{aligned} \tag{16}$$

where Ω is the feasible design space defined by the variable ranges in Table 4 and $g_i(\mathbf{x})$ represents the set of equality and inequality constraints. The decision vector ($\mathbf{x}$) includes powertrain sizing parameters, ADU catalyst design variables, and placeholder control-strategy parameters:

$$\begin{aligned} \mathbf{x} = [P_{FC}^{\max}, P_{mot}^{\max}, Q_{bat}, \rho_1^{cat\_type}, \rho_2^{cat\_type}, \rho_3^{cat\_type}, \rho_4^{cat\_type}, \rho_5^{cat\_type}, \\ \rho_1^{cat\_prop}, \rho_2^{cat\_prop}, \rho_3^{cat\_prop}, \rho_4^{cat\_prop}, \rho_5^{cat\_prop}, \lambda_0, k_p, k_d]^T \end{aligned} \tag{17}$$

Table 4. Definition of multi-objective co-optimization problem.

Parameters	**Definition**
Objective functions	$\min \left[\mathbf{f} = \left(J_{fuel}, J_{SOH}, J_{cost}\right)\right]$
	$J_{fuel} = \int_{t_0}^{t_f} \dot{m}_{H_2}(t)\, dt$
	$J_{SOH} = \Delta SOH_{FC} + \Delta SOH_{bat}$
	$J_{cost} = \sum_{i=\text{FC,mot,bat,...}} C_{cost}^i$
Decision variables	$150\ \text{kW} \leq P_{FC}^{max} \leq 350\ \text{kW}$
	$50\ \text{kW} \leq P_{mot}^{max} \leq 150\ \text{kW}$
	$50\ \text{Ah} \leq Q_{bat} \leq 100\ \text{Ah}$
	$\rho_i^{cat\_type} \in \{A, B, \ldots, J\}, \quad i \in \{I, II, III, IV, V\}$
	$0 \leq \rho_i^{cat\_prop} \leq 1, \quad i \in \{I, II, III, IV, V\}$
	$\lambda_0 \in [-3, 3]$
	$k_p^{\lambda} \in [0, 20]$
	$k_i^{\lambda} \in [0, 2]$
Constraints	$\sum_{i=I}^{V} \rho_i^{cat} = 1$
	$P_{ICE}^{min} \leq P_{ICE} \leq P_{ICE}^{max}$
	$P_{DM}^{min} \leq P_{DM} \leq P_{DM}^{max}$
	$P_{bat}^{min} \leq P_{bat} \leq P_{bat}^{max}$
	$SOC_{min} \leq SOC \leq SOC_{max}$

The high dimensionality of the decision space, the non-linear and tightly coupled nature of the system dynamics, and the computationally expensive nature of full-mission simulations render direct optimization intractable. Traditional surrogate-assisted multi-objective evolutionary algorithms typically follow a decoupled paradigm: a data-driven surrogate model is trained to approximate the fitness landscape, and a gradient-free multi-objective evolutionary algorithm, such as NSGA-II, uses this surrogate for fitness evaluations. This approach suffers from two critical limitations in the context of flying vehicle design:

1. Physical inconsistency: Black-box surrogates have no intrinsic knowledge of the system's governing laws. In data-sparse regions, their predictions can violate funda-

mental principles like energy conservation, leading the optimizer towards physically unrealizable local optimality.

2. Inefficient search: Gradient-free evolutionary algorithms operate via stochastic operators (crossover and mutation). While robust for global exploration, this "blind" search is inefficient for the fine-tuning of promising solutions. They lack a sense of direction and do not exploit the local topology of the fitness landscape to accelerate convergence.

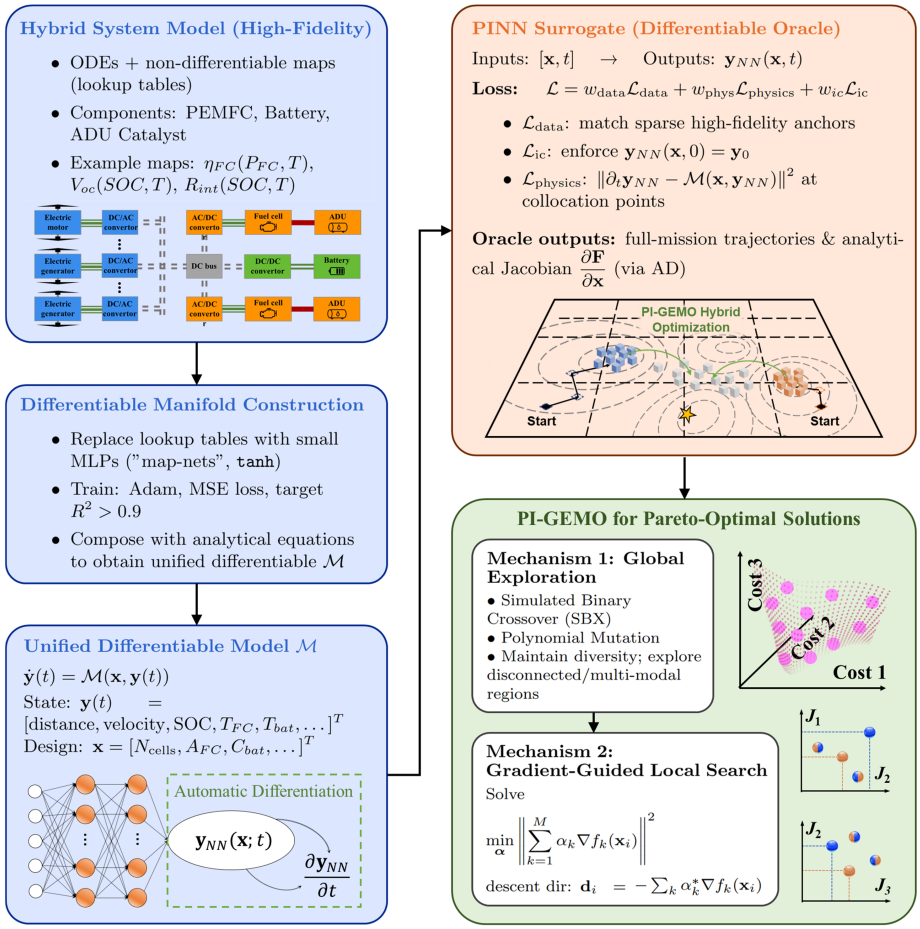


Figure 5. The overall conceptual design of the PI-GEMO framework for the co-optimization of the ammonia-powered flying vehicle propulsion system and control parameters, where the start refers to the optimal solution.

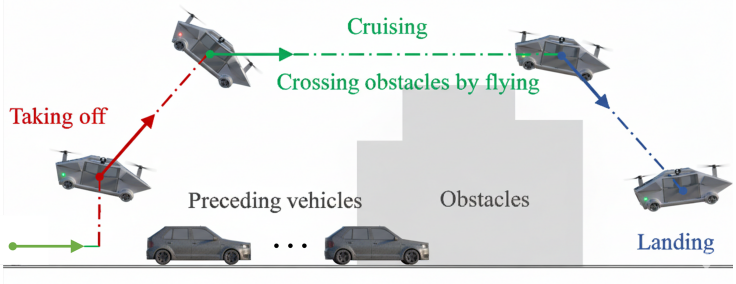


Figure 6. Scenario setting of a typical 2condition for HEVTOLs.

3.2. The Physics-Informed Differentiable Surrogate Model

The fidelity of the PI-GEMO framework is fundamentally tied to the physical model embedded within the PINN. For this study, the model is defined by the governing equations detailed in the modeling methodology. We explicitly state that this model does not account for certain second-order or long-term degradation effects, and its scope is therefore focused

on the co-design of component sizing and control based on the dominant, well-understood physics of the system [40]. The cornerstone of the PI-GEMO framework is a surrogate model that is not only computationally efficient but also physically consistent and fully differentiable. The construction is a two-stage process: first, transforming the entire hybrid system model into a continuously differentiable manifold and, second, training a PINN to learn the system's dynamic response on this manifold. This approach ensures that the powerful tools of gradient-based optimization can be fully leveraged.

3.2.1. Construction of the Differentiable System Manifold

The high-fidelity model developed in Section 2 is a hybrid system, comprising a mix of analytical ordinary differential equations (ODEs) and non-differentiable, map-based components (e.g., lookup tables for PEMFC efficiency, battery SOC-dependent internal resistance, etc.). To enable the use of gradient-based methods, it is imperative to first construct a fully differentiable representation of the entire system. This is accomplished by replacing each discrete lookup table and other non-analytical components with a continuous and smooth function approximator. We utilize small, fully-connected neural networks for this task, which we term "map-nets".

Each map-net is an independent, shallow multi-layer perceptron (MLP) with hyperbolic tangent ('tanh') activation functions in its hidden layers. The 'tanh' function is specifically chosen over alternatives like the rectified linear unit 'ReLU' because it is infinitely differentiable (C^∞-continuous). This property is crucial for the numerical stability and accuracy of the higher-order gradient computations required by both the PINN training and the subsequent multi-objective optimization steps, avoiding issues like vanishing or exploding gradients. Each map-net, parameterized by weights and biases ($\boldsymbol{\phi}$), is trained via standard backpropagation using the Adam optimizer to minimize the mean squared error (MSE) against the original map data until a high degree of fidelity is achieved with an indicator of R^2. For a generic map where a set of outputs ($\mathbf{z}$) is a function of inputs ($\mathbf{u}$), the corresponding map-net ($f_{NN}(\mathbf{u};\boldsymbol{\phi})$) is trained by solving the unconstrained optimization problem:

$$\min_{\boldsymbol{\phi}} \mathcal{L}_{map}(\boldsymbol{\phi}) = \min_{\boldsymbol{\phi}} \frac{1}{N_{map}} \sum_{i=1}^{N_{map}} \| f_{NN}(\mathbf{u}_i;\boldsymbol{\phi}) - \mathbf{z}_i \|_2^2 \tag{18}$$

This process is meticulously applied to all map-based components, yielding a library of differentiable sub-models that seamlessly integrate into the larger system dynamics.

3.2.2. Physics-Informed Surrogate Learning

The PINN, denoted by $\mathbf{y}_{NN}(\mathbf{x}, t; \boldsymbol{\theta})$, is a deep neural network with parameters ($\boldsymbol{\theta}$) designed to learn the solution trajectory of the system's ODEs for a given design vector ($\mathbf{x}$) over the entire mission time horizon ($t \in [0, T]$). Its architecture is a deep MLP, again using 'tanh' activation functions to ensure high-order differentiability. The network takes the concatenated vector ($[\mathbf{x}, t]$) as input and outputs the predicted state vector ($\mathbf{y}_{NN}$) at that time.

The PINN is trained by minimizing a composite loss function that balances fidelity to known data with adherence to the governing physical laws:

$$\mathcal{L}(\boldsymbol{\theta}) = w_{data}\mathcal{L}_{\text{data}}(\boldsymbol{\theta}) + w_{phys}\mathcal{L}_{\text{physics}}(\boldsymbol{\theta}) + w_{ic}\mathcal{L}_{\text{ic}}(\boldsymbol{\theta}) \tag{19}$$

where w_{data}, w_{phys}, and w_{ic} are hyperparameters that weight the contribution of each loss term. The data loss term ($\mathcal{L}_{\text{data}}$) anchors the model to reality using a sparse set of N_s

high-fidelity simulation results. This ensures that the PINN learns to match the true system trajectories at specific anchor points, preventing divergence.

$$\mathcal{L}_{\text{data}} = \frac{1}{N_s} \sum_{i=1}^{N_s} \|\mathbf{y}_{NN}(\mathbf{x}_i, t_i; \boldsymbol{\theta}) - \mathbf{y}_{\text{sim}}(\mathbf{x}_i, t_i)\|_2^2 \tag{20}$$

The initial condition loss term ($\mathcal{L}_{\text{ic}}$) specifically enforces the requirement that the PINN solution satisfy the known initial state of the system ($\mathbf{y}_0$) for any given design ($\mathbf{x}$).

$$\mathcal{L}_{\text{ic}} = \frac{1}{N_{ic}} \sum_{j=1}^{N_{ic}} \left\|\mathbf{y}_{NN}(\mathbf{x}_j, t = 0; \boldsymbol{\theta}) - \mathbf{y}_0\right\|_2^2 \tag{21}$$

In addition, the physics loss ($\mathcal{L}_{\text{physics}}$) enforces the system's governing dynamics by penalizing the residual of the ODEs defined by our differentiable manifold ($\mathcal{M}$). This residual is computed at a large number of N_c random collocation points $(\mathbf{x}_j, t_j)$ sampled from the design and time domains. The time derivative of the PINN's output ($\frac{\partial \mathbf{y}_{NN}}{\partial t}$) is calculated analytically using automatic differentiation (AD), not numerical approximation, which is key to its accuracy.

$$\mathcal{L}_{\text{physics}} = \frac{1}{N_c} \sum_{j=1}^{N_c} \left\| \frac{\partial \mathbf{y}_{NN}(\mathbf{x}_j, t_j; \boldsymbol{\theta})}{\partial t} - \mathcal{M}(\mathbf{x}_j, \mathbf{y}_{NN}(\mathbf{x}_j, t_j; \boldsymbol{\theta})) \right\|_2^2 \tag{22}$$

By minimizing this composite loss, the PINN is forced to learn a continuous function that is not just close to the data points but is also a valid solution to the system's governing equations across the entire design time space. The trained surrogate provides near-instantaneous predictions of the full mission trajectory and, crucially, the analytical Jacobian of the objective functions ($\mathbf{F}(\mathbf{x})$) with respect to the decision variables ($\mathbf{x}$). For a set of M objectives ($f_1, \ldots, f_M$) and D design variables ($x_1, \ldots, x_D$), the Jacobian matrix ($\mathbf{J}$) at a point ($\mathbf{x}_i$) is computed via AD through the network:

$$\mathbf{J}_i = \frac{\partial \mathbf{F}(\mathbf{x}_i)}{\partial \mathbf{x}_i} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \cdots & \frac{\partial f_1}{\partial x_D} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \cdots & \frac{\partial f_2}{\partial x_D} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_M}{\partial x_1} & \frac{\partial f_M}{\partial x_2} & \cdots & \frac{\partial f_M}{\partial x_D} \end{bmatrix}_{\mathbf{x}=\mathbf{x}_i} \tag{23}$$

3.3. The PI-GEMO Hybrid Optimization Algorithm

PI-GEMO integrates the gradient information from the PINN directly into its search loop, creating a hybrid algorithm that synergizes two complementary search mechanisms: gradient-free global exploration and gradient-guided local search. This hybrid nature allows it to balance the need for broad exploration of the design space (exploration) with rapid convergence to high-precision solutions on the true Pareto front (exploitation).

3.3.1. Mechanism 1: Gradient-Free Global Exploration

To ensure the algorithm thoroughly explores the design space and avoids premature convergence to local or non-convex regions of the Pareto front, PI-GEMO employs robust, gradient-free genetic operators. This mechanism is responsible for maintaining population diversity and discovering new, potentially disconnected regions of the Pareto front. Drawing inspiration from natural evolution, we use two well-established operators:

- Simulated binary crossover (SBX): This operator takes two parent solutions ($\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$) and creates two offspring ($\mathbf{c}^{(1)}$ and $\mathbf{c}^{(2)}$) by simulating the behavior of a single-

point crossover on binary strings. The process for each variable (j) is governed by a spread factor β_j derived from a random number ($u_j \in [0,1]$) and a distribution index (η_c). SBX is adept at combining features from good solutions to generate potentially superior new ones within the hyper-rectangle defined by the parents. The offspring are generated as follows:

$$\begin{cases} \mathbf{c}_j^{(1)} & = 0.5\left[(1-\beta_j)\mathbf{x}_j^{(1)} + (1+\beta_j)\mathbf{x}_j^{(2)}\right] \\ \mathbf{c}_j^{(2)} & = 0.5\left[(1+\beta_j)\mathbf{x}_j^{(1)} + (1-\beta_j)\mathbf{x}_j^{(2)}\right] \end{cases} \tag{24}$$

- Polynomial mutation: This operator introduces small, localized perturbations to a solution vector ($\mathbf{x}$), mimicking the effect of random mutation in nature. For each variable ($\mathbf{x}_j$), a perturbation (δ_j) is calculated based on a random number and a distribution index (η_m). The mutated variable ($\mathbf{x}'_j$) is given as follows, where $\mathbf{x}_j^U$ and $\mathbf{x}_j^L$ are the upper and lower bounds for the variable. This operator is crucial for the fine-tuning of solutions and the exploration of the immediate neighborhood of existing points on the Pareto front.

$$\mathbf{x}'_j = \mathbf{x}_j + \delta_j(\mathbf{x}_j^U - \mathbf{x}_j^L) \tag{25}$$

These stochastic operators require no gradient information, making them effective at navigating the global and potentially deceptive and multi-modal landscape of the multi-objective problem.

3.3.2. Mechanism 2: Gradient-Guided Local Search

For a given non-dominated solution ($\mathbf{x}_i$) in the current population, this mechanism uses the gradient information from the PINN to identify the most efficient path towards the Pareto front. The theoretical basis for this search lies in the Karush–Kuhn–Tucker (KKT) first-order necessary conditions for multi-objective optimization. A point ($\mathbf{x}_i$) is Pareto-optimal only if there exists no single descent direction ($\mathbf{d}$) for which all objectives can be improved simultaneously. This implies that the negative gradients' convex cone must not have a common intersection or, equivalently, that the convex hull of the objective gradients ($\{\nabla f_1(\mathbf{x}_i), \ldots, \nabla f_M(\mathbf{x}_i)\}$) must contain the zero vector. If it does not, a common descent direction exists. We find this direction by solving the following quadratic programming (QP) sub-problem, which seeks the minimum-norm vector in the convex hull of the objective gradients [41]:

$$\begin{aligned} \min_{\boldsymbol{\alpha}} \quad & \left\| \sum_{k=1}^{M} \alpha_k \nabla f_k(\mathbf{x}_i) \right\|_2^2 \\ \text{s.t.} \quad & \sum_{k=1}^{M} \alpha_k = 1, \quad \alpha_k \geq 0 \quad \forall k \in \{1, \ldots, M\} \end{aligned} \tag{26}$$

where M is the number of objectives and $\nabla f_k(\mathbf{x}_i)$ is the k-th row of the Jacobian matrix ($\mathbf{J}_i$). The solution to this QP problem ($\boldsymbol{\alpha}^* = [\alpha_1^*, \ldots, \alpha_M^*]^T$) provides the optimal weights to combine the gradients. The resulting vector ($\mathbf{d}_i$) is the steepest common descent direction:

$$\mathbf{d}_i = -\sum_{k=1}^{M} \alpha_k^* \nabla f_k(\mathbf{x}_i) \tag{27}$$

Once the descent direction ($\mathbf{d}_i$) is found a backtracking line search is performed to find an appropriate step size ($\beta > 0$). The search starts with an initial step size of β_0 and iteratively reduces it ($\beta \leftarrow \rho\beta$ with $\rho \in (0,1)$) until the new point ($\mathbf{x'}_i = \mathbf{x}_i + \beta\mathbf{d}_i$) satisfies

the multi-objective Armijo condition. This condition ensures a sufficient decrease in the directionally weighted objective function without unduly worsening others:

$$\max_{k\in\{1,\dots,M\}}\{f_k(\mathbf{x}_i+\beta\mathbf{d}_i)-f_k(\mathbf{x}_i)\}\leq c\beta\max_{k\in\{1,\dots,M\}}\left\{\nabla f_k(\mathbf{x}_i)^T\mathbf{d}_i\right\} \tag{28}$$

where $c\in(0,1)$ is a control parameter. This gradient-guided step allows the solution to move rapidly and efficiently towards the true Pareto front, dramatically accelerating the convergence of the algorithm.

The complete workflow of the PI-GEMO algorithm, which integrates the PINN surrogate with the hybrid search mechanisms, is rigorously defined in Algorithm 1. In each generation, the algorithm begins by evaluating the current population using the PINN. It then performs non-dominated sorting to identify the current Pareto front. For the elite solutions on this front, the gradient-guided local search (Mechanism 2) is activated. The gradients are queried from the PINN; the QP sub-problem (Equation (26)) is solved for each solution to find the optimal descent direction; and a line search is performed to produce a new, improved set of elite solutions. Concurrently, the rest of the population undergoes gradient-free global exploration (Mechanism 1), where SBX and polynomial mutation operators are applied to generate a diverse set of offspring. Finally, the improved elites and the exploratory offspring are combined to form the population for the next generation. This cycle of evaluation, gradient-based refinement, and global exploration continues until a termination criterion, such as a maximum number of generations or a convergence metric on the Pareto front, is met.

Algorithm 1 The PI-GEMO Algorithm.

Require: MOP, Population size N, Max generations G_{max}, Local search frequency F_{ls}
Ensure: Final Pareto-optimal set P_{final}

Create differentiable system representation $\mathcal{M}$ by training map-nets on all lookup tables.
Generate a sparse high-fidelity dataset $\mathcal{D}_{sim}$ using Latin hypercube sampling.
Train the PINN surrogate $\mathbf{y}_{NN}(\mathbf{x},t;\boldsymbol{\theta})$ by minimizing Equation (19) using $\mathcal{M}$ and $\mathcal{D}_{sim}$.
Initialize a random population P_0 of size N.
Evaluate $\mathbf{F}(\mathbf{x})$ for all $\mathbf{x}\in P_0$ using the trained PINN.
for $g=1$ to G_{max} **do**
    Generate offspring population Q_g from P_g using SBX and polynomial mutation.
    Initialize $R_g=\emptyset$.
    if $g \pmod{F_{ls}} == 0$ **then**
        Determine the non-dominated front P'_g from P_g.
        for each individual $\mathbf{x}_i\in P'_g$ **do**
            Compute Jacobian $\mathbf{J}_i=\frac{\partial\mathbf{F}(\mathbf{x}_i)}{\partial\mathbf{x}_i}$ using the PINN and AD.
            Solve Equation (26) to find weights $\boldsymbol{\alpha}^*$.
            Compute descent direction $\mathbf{d}_i=-\mathbf{J}_i^T\boldsymbol{\alpha}^*$.
            Perform line search along $\mathbf{d}_i$ to find step size β.
            Create improved individual $\mathbf{x'}_i=\mathbf{x}_i+\beta\mathbf{d}_i$.
            $R_g\leftarrow R_g\cup\{\mathbf{x'}_i\}$.
        end for
    end if
    Combine populations: $C_g=P_g\cup Q_g\cup R_g$.
    Evaluate all new individuals in C_g using the PINN.
    Select the next generation P_{g+1} of size N from C_g using the standard NSGA-II selection mechanism (non-dominated sorting and crowding distance).
end for
return The non-dominated set from the final population $P_{G_{max}}$.

4. Results and Discussions

This section presents a comprehensive evaluation of the proposed PI-GEMO framework. The analysis is structured to first define the case study and experimental setup, then validate the foundational PINN surrogate model and subsequently compare the optimization process and final results against a state-of-the-art NSGA-III benchmark.

4.1. Case Study and Experimental Setup

The optimization and validation processes are conducted for representative flying vehicle missions, combining both flying and road-driving conditions, as depicted in Figure 7. All optimizations are performed with a population size of 100 and run for a maximum of 150 generations. The cost of flying vehicle development is summarized in Table 5. The proposed PI-GEMO framework is benchmarked against a standard NSGA-III algorithm. To ensure a fair comparison, ablation experiments were conducted to examine the differences between combining PINN with the standard NSGA-III (PI-NSGA) and integrating the PINN with the proposed GEMO framework (PI-GEMO). The benchmark algorithm, PI-NSGA, uses the exact same trained PINN for its fitness evaluations but lacks the gradient-guided local search mechanism, relying solely on stochastic crossover and mutation operators.

Table 5. The cost of flying vehicle development [42].

Components	Cost Price C^i_{size} (USD)
Body and chassis	85,000
Fuel cell and control system	175 $P^{\text{FC}}_{\text{max}}$
Motor and control unit	21.775 $P^{\text{mot}}_{\text{max}} + 425$
Battery and control unit	200 E_{bat}

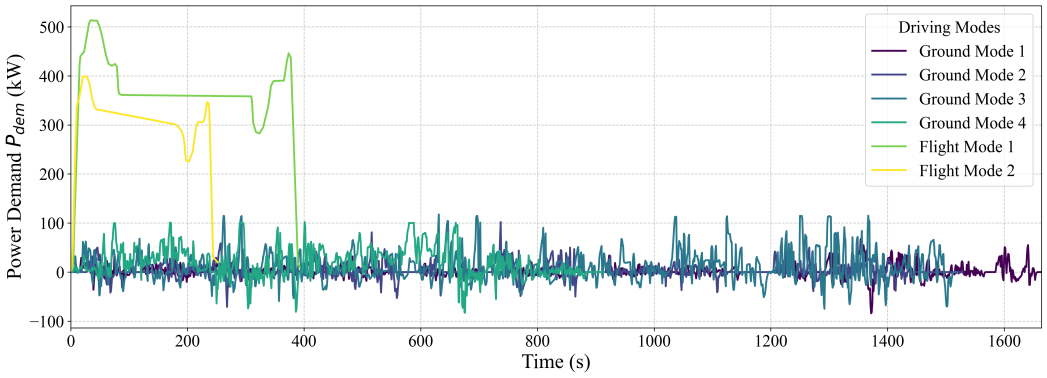


Figure 7. The employed drive cycles integrating both terrestrial and aerial modes.The defined flight-mission profile for the flying vehicle logistics case study, including takeoff, cruise, hover, and landing phases, used to evaluate the objective functions.

4.2. Validation of the Physics-Informed Surrogate Model and Gradient-Enhanced Pareto Optimization

The credibility of the entire optimization framework hinges on the accuracy of the PINN surrogate. To validate its performance, a separate set randomly sampled from the design space and not used during training was generated using the high-fidelity simulator. Figure 8 compares the PINN's predictions for the two primary objective functions against the true values from the simulator. The coefficient of determination (R^2) was found to be 0.971 for hydrogen consumption prediction and 0.953 for total SOH loss. This confirms that the PINN, by leveraging both data and the underlying physical equations, successfully learned the complex, non-linear mapping from the high-dimensional design space to the performance space. This high accuracy validates its use as a reliable and computationally efficient surrogate for the subsequent optimization algorithms.

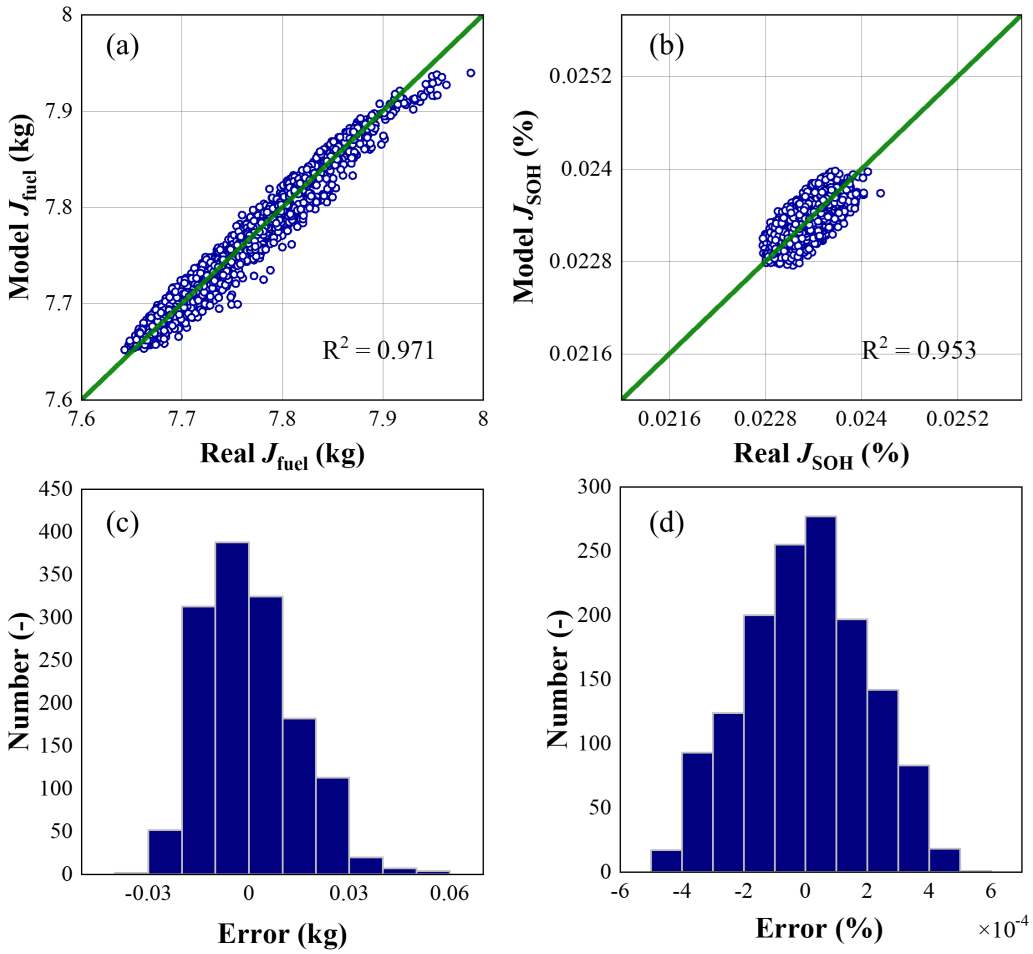


Figure 8. Validation of the trained PINN surrogate model: (**a**) predicted vs. real fuel consumption; (**b**) predicted vs. real SOH loss; (**c**) statistical histogram of model deviation for fuel consumption; (**d**) statistical histogram of model deviation for SOH loss.

With the surrogate validated, we now compare the optimization efficacy of the proposed PI-GEMO against the benchmark NSGA-III during the evolutionary process. The final non-dominated fronts obtained after 150 generations are shown in Figure 9. The PI-GEMO front clearly and substantially dominates the final NSGA-III front, extending further towards both the maximum endurance and minimum mass objectives. To understand the physical implications of these results, we select three representative solutions from each front for detailed analysis: the design that maximizes flight endurance ('Max Endurance'), the design that minimizes takeoff mass ('Min Mass'), and a balanced 'Trade-off' solution. Figure 9 illustrates the evolution of the Pareto fronts for both algorithms at different generational milestones. It is visually evident that the PI-GEMO framework achieves a superior distribution of solutions from a very early stage. By the 20th generation, PI-GEMO has already identified a front that significantly dominates the one found by NSGA-III at the same stage. This demonstrates the immediate impact of the gradient-guided search in rapidly identifying promising regions of the design space.

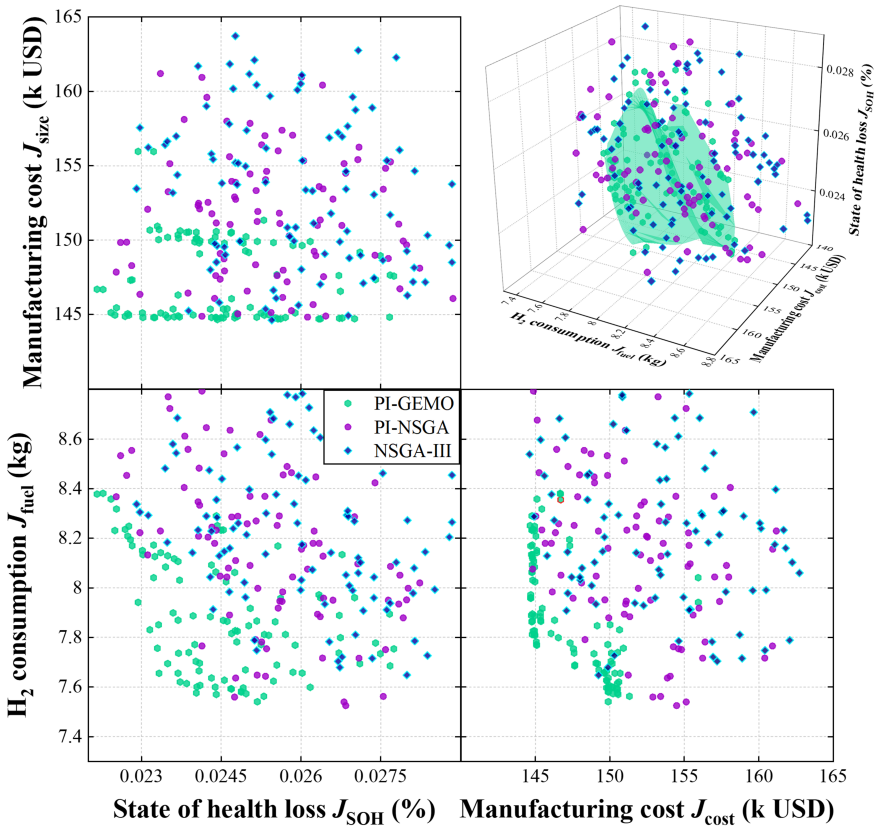


Figure 9. Comparison of the final optimization solutions obtained by PI-GEMO, PI-NSGA, and the NSGA-III benchmark after 60 generations.

To quantitatively assess the convergence and quality of the Pareto fronts, the inverted generational distance (IGD) metric is employed. The IGD measures the average Euclidean distance from a set of uniformly distributed reference points on the true (or best-known) Pareto front to the solutions found by the algorithm. A lower IGD value signifies that the algorithm's solutions are both closer to the front and more evenly distributed along it. Figure 10 plots the IGD values for both algorithms over the 80 generations. The PI-GEMO algorithm exhibits a much steeper initial drop in IGD, reaching a low value within the first 30 generations, after which it continues to make steady fine-tuning improvements. In contrast, the NSGA-II algorithm shows a much slower, more gradual decrease in IGD and plateaus at a significantly higher value. This quantitative result confirms the visual evidence from Figure 9 and powerfully demonstrates the effectiveness of the gradient-guided local search mechanism.

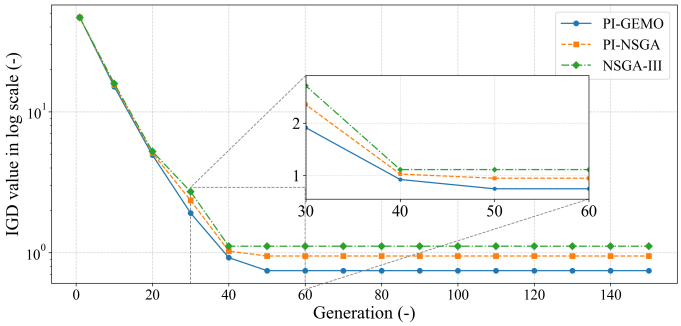


Figure 10. Convergence dynamics of the inverted generational distance metric for PI-GEMO, PI-NSGA, and the NSGA-III benchmark.

4.3. Comparisons of Optimization Results and Sensitivity Analysis

The trade-offs and the dominance of the PI-GEMO results are visualized in Figures 11 and 12, which compare the normalized performance of the 'Trade-off' solutions from both algorithms. The larger area covered by the PI-GEMO solution graphically illustrates its superiority in achieving a better balance between the conflicting objectives. The design parameters and corresponding objective values for these selected points are presented in Table 6. The PI-GEMO algorithm converges to a distinctly different powertrain architecture. It selects a smaller fuel cell ($P_{\max}^{\mathrm{FC}}$ = 249.2 kW) compared to both PI-NSGA (255.7 kW) and NSGA-III (271.4 kW). This reduction in fuel-cell size is a primary driver for the lower system cost (J_{size}) and mass. To compensate for the lower peak power from the fuel cell, the system relies on a well-sized battery pack (Q_{bat} = 84.5 Ah) and motor ($P_{\max}^{\mathrm{mot}}$ = 91.9 kW) to handle transient power demands, reflecting a more optimized power-sharing strategy.

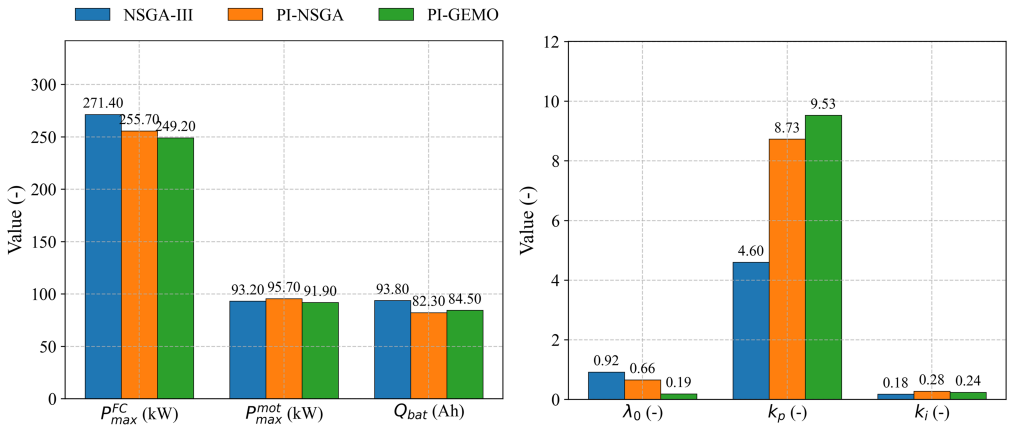


Figure 11. Normalized performance comparison of the trade-off solutions found by PI-GEMO, PI-NSGA, and NSGA-III.

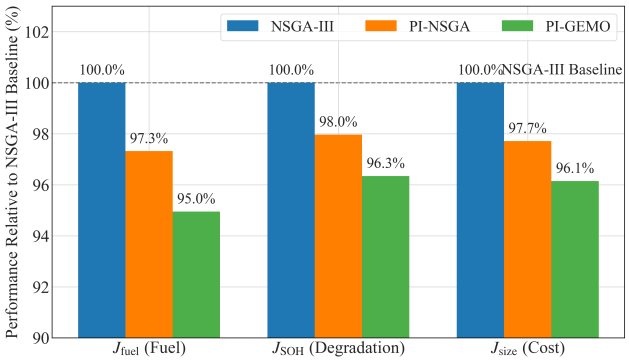


Figure 12. Normalized performance comparison of the trade-off solutions found by PI-GEMO, PI-NSGA, and NSGA-III.

The control strategy parameters show a clear trend. The PI-GEMO solution features a significantly lower initial power-sharing factor (λ_0 = 0.19) compared to NSGA-III (0.92). This indicates a control strategy that is inherently more reliant on the battery for initial power response, preserving the fuel cell for more stable, high-efficiency operation. This is coupled with a more aggressive proportional gain (k_p^{λ} = 9.53), suggesting a controller that can react more swiftly to changes in power demand, a strategy likely enabled by the accurate gradient information. The optimization of the catalyst composition provides fascinating insights. The baseline NSGA-III solution utilizes a mix including catalyst 'B'. In contrast, both gradient-enhanced methods, i.e., PI-NSGA and PI-GEMO, discard catalyst 'B' in favor of a combination involving higher-activity catalysts like 'D', 'E', 'G', and 'I'. The

PI-GEMO solution further refines this by settling on a [D, E, E, E, G] vector, suggesting that a higher proportion of catalyst 'E' is optimal for this balanced design. Moreover, the PI-GEMO solution exhibits a more uniform catalyst proportion vector, with all components near 20%. This suggests that the gradient information guided the search away from uneven distributions towards a more homogenized and synergistic catalyst blend, maximizing overall efficiency and longevity.

Table 6. Performance and design parameters of the 'Trade-off' solutions found by the benchmark and proposed optimization algorithms.

Parameter	NSGA-III	PI-NSGA	PI-GEMO
$P_{\max}^{FC}$ (kW)	271.4	255.7	249.2
$P_{\max}^{mot}$ (kW)	93.2	95.7	91.9
Q_{bat} (Ah)	93.8	82.3	84.5
λ_0 (-)	0.92	0.66	0.19
k_p^{λ} (-)	4.60	8.73	9.53
k_i^{λ} (-)	0.18	0.28	0.24
Catalysts Type Vector (-)	[B, D, E, E, G]	[D, E, E, G, I]	[D, E, E, E, G]
Catalysts Proportion Vector (%)	[18.81, 20.69, 19.19, 21.76, 19.55]	[19.34, 20.30, 19.43, 20.71, 20.22]	[20.16, 19.60, 19.94, 19.89, 20.41]
J_{fuel} (kg)	7.903	7.691	7.504
J_{SOH} (%)	0.0246	0.0241	0.0237
J_{size} (k USD)	150.917	147.464	145.105

To further understand the design space and validate the optimization results, a local sensitivity analysis was performed. By leveraging the auto-differentiable nature of the PINN surrogate, we can instantaneously extract the gradients of the objective functions with respect to each normalized design variable from the PINN's Jacobian matrix. This analysis was centered around the PI-GEMO 'Trade-off' solution, and the results are visualized in Figure 13.

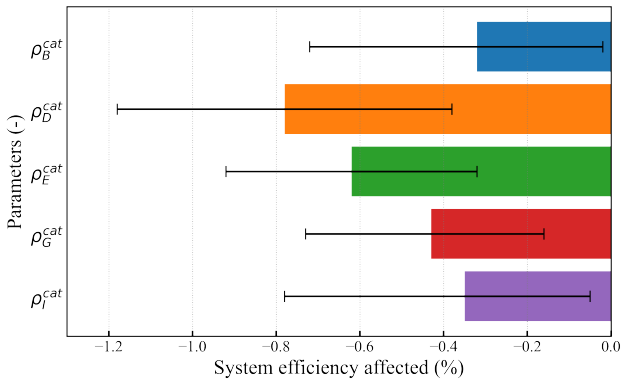


Figure 13. Sensitivity analysis of catalyst design for system efficiency.

The analysis reveals that different parameters exhibit vastly different levels of influence on the system's performance and cost. For system efficiency (inversely related to J_{fuel}), the analysis highlights the catalyst ratios as highly influential. As shown in Figure 14 (adapted from the general sensitivity analysis of Figure 13), the proportion of catalyst type 'D' exhibits the highest positive sensitivity, meaning an increase in its proportion yields the most significant gain in efficiency. This finding strongly corroborates the design choice made by the PI-GEMO algorithm, which selected a catalyst combination rich in types 'D' and 'E'. This confirms that our conclusions are not arbitrary but are rooted in the physical sensitivities of the system.

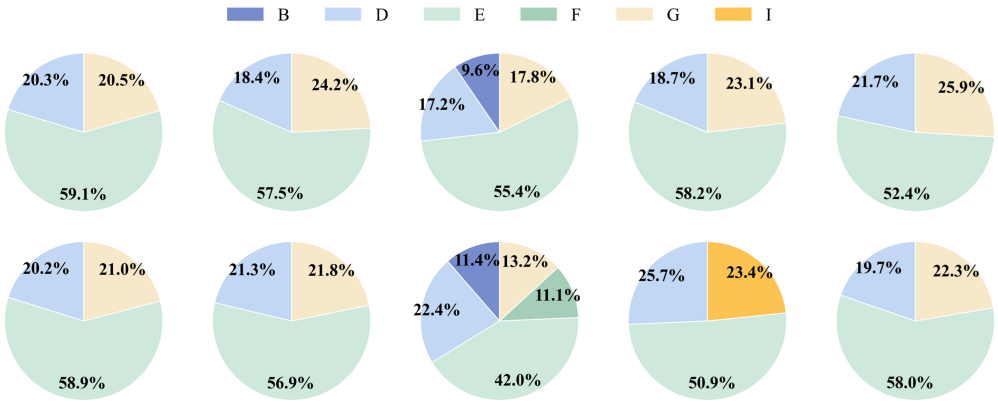


Figure 14. Gradients of the 10 selected ADU design results of the PI-GEMO-optimized solutions in the 60th generation.

5. Conclusions

This paper addresses the formidable challenge of co-designing ammonia-powered fuel-cell flying vehicles, a problem characterized by tight multi-physics couplings, a high-dimensional design space, and computationally intensive simulations. To this end, a novel optimization paradigm, i.e., PI-GEMO, was proposed and validated. The PI-GEMO framework uniquely integrates a differentiable PINN surrogate model into a hybrid evolutionary algorithm, synergizing the global exploration capabilities of genetic algorithms with the rapid convergence of gradient-based local search. Key contributions and findings include the follow:

- A high-fidelity, multi-physics dynamic model of an ammonia-powered fuel-cell flying vehicle was developed, holistically capturing the critical couplings between the physical, thermal, electrical, and aerodynamic domains. This model formed the basis for the training of a differentiable PINN surrogate, which not only accelerates fitness evaluations but, by incorporating the system's governing equations into its loss function, critically ensures the physical plausibility and accuracy of the optimization results.
- The proposed PI-GEMO framework introduces a new class of hybrid intelligent optimization. It combines traditional, gradient-free genetic operators for robust global exploration with a novel gradient-guided local search mechanism. This mechanism leverages analytical gradients extracted from the trained PINN via automatic differentiation to efficiently propel promising solutions towards the true Pareto-optimal front. The superiority of this hybrid approach was demonstrated through a significantly faster convergence rate and a lower final IGD value compared to the benchmark NSGA-III algorithm.
- A comprehensive case study demonstrated the practical efficacy of the PI-GEMO framework. The co-optimization of ADU catalyst composition, powertrain component sizing, and energy management control parameters yielded a set of Pareto-optimal solutions that substantially dominated those found by conventional methods. Specifically, the PI-GEMO-derived 'Trade-off' design achieved a simultaneous reductions in hydrogen consumption of 5.1%, power source degradation SOH of 3.7%, and total system cost of 3.9% when compared to the solution from a standard NSGA-III optimization, providing critical insights into the synergistic design of next-generation aerial propulsion systems.
- While this study focused on a hexacopter architecture, the proposed PI-GEMO methodology is highly generalizable. The framework's core—the synergistic integration of a physics-informed differentiable surrogate with a hybrid gradient-enhanced

algorithm—is fundamentally model-agnostic. It can be readily adapted to other complex engineering systems, such as eVTOLs with different rotor counts, fixed-wing hybrid aircraft, or even terrestrial vehicles, by simply replacing the underlying set of governing equations within the PINN's physics loss function. This scalability makes PI-GEMO a powerful and versatile tool for a wide range of multi-domain co-design problems. Future work will also focus on incorporating higher-fidelity transient models for both the ADU thermal dynamics and battery electrochemistry to further refine the design of the control strategy.

Author Contributions: Methodology, Y.B. and N.L.; software, Y.B. and N.L.; validation, Y.B.; data curation, Y.B.; writing—original draft, Y.B.; writing—review and editing, N.L. and H.Z.; supervision, C.C.; funding acquisition, C.C. and H.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Natural Science Foundation of China (Grant No. 52302410).

Data Availability Statement: The original contributions presented in the study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: The assistance of the research fellows at Beijing Tsintelink Technology Co., Ltd., in model development and coding is gratefully acknowledged, and the code will be made available upon request for research purposes.

Conflicts of Interest: The authors declare that this study received support from Beijing Tsintelink Technology Co., Ltd. Beijing Tsintelink Technology Co., Ltd., was not involved in the study design; the collection, analysis, or interpretation of data; the writing of this article; or the decision to submit it for publication.

References

1. Lan, R.; Irvine, J.T.S.; Tao, S. Ammonia and related chemicals as potential indirect hydrogen storage materials. *Int. J. Hydrogen Energy* **2012**, *37*, 1482–1494. [CrossRef]
2. Milićević, S.; Blagojević, I.; Milojević, S.; Bukvić, M.; Stojanović, B. Numerical Analysis of Optimal Hybridization in Parallel Hybrid Electric Powertrains for Tracked Vehicles. *Energies* **2024**, *17*, 3531. [CrossRef]
3. Hepperle, M. Electric Flight—Potential and Limitations. In Proceedings of the 30th International Congress of the Aeronautical Sciences (ICAS 2012), Brisbane, Australia, 23–28 September 2012.
4. Zhang, H.; Lei, N.; Chen, B.; Li, B.; Li, R.; Wang, Z. Modeling and control system optimization for electrified vehicles: A data-driven approach. *Energy* **2024**, *310*, 133196. [CrossRef]
5. Koffler, J.; Fansler, T.D. Hydrogen Fuel for Transportation: A Review. *SAE Int. J. Fuels Lubr.* **2021**, *14*, 1–22.
6. Meng, X.; Sun, C.; Mei, J.; Tang, X.; Hasanien, H.M.; Jiang, J.; Fan, F.; Song, K. Fuel cell life prediction considering the recovery phenomenon of reversible voltage loss. *J. Power Sources* **2025**, *625*, 235634. [CrossRef]
7. Meng, X.; Mei, J.; Tang, X.; Jiang, J.; Sun, C.; Song, K. The Degradation Prediction of Proton Exchange Membrane Fuel Cell Performance Based on a Transformer Model. *Energies* **2024**, *17*, 3050. [CrossRef]
8. Ahluwalia, R.K.; Hua, T.Q.; Peng, J.K.; Lasher, S.; McKenney, K.; Sinha, J.; Weisberg, A. On-board and off-board performance of hydrogen storage systems for light-duty vehicles. *Int. J. Hydrogen Energy* **2010**, *35*, 12339–12353. [CrossRef]
9. Valera-Medina, A.; Xiao, H.; Owen-Jones, M.; David, W.I.F.; Bowen, P.J. Ammonia for power. *Prog. Energy Combust. Sci.* **2018**, *69*, 63–102. [CrossRef]
10. Zhang, H.; Lei, N.; Peng, W.; Li, B.; Lv, S.; Chen, B.; Wang, Z. Bi-Level Transfer Learning for Lifelong-Intelligent Energy Management of Electric Vehicles. *IEEE Trans. Intell. Transp. Syst.* **2025**, 1–14. [CrossRef]
11. International Energy Agency. *The Role of Low-Carbon Fuels in the Clean Energy Transitions of the Power Sector*; Technical Report; IEA: Paris, France, 2021.
12. Zhou, J.; Shu, X.; Zhang, J.; Yi, F.; Jia, C.; Zhang, C.; Kong, X.; Zhang, J.; Wu, G. A deep learning method based on CNN-BiGRU and attention mechanism for proton exchange membrane fuel cell performance degradation prediction. *Int. J. Hydrogen Energy* **2024**, *94*, 394–405. [CrossRef]
13. Wang, Y.; Fan, F.; Wang, Y.; Wang, K.; Jiang, J.; Sun, C.; Xue, R.; Song, K. Convective Heat Loss Prediction Using the Concept of Effective Wind Speed for Dynamic Line Rating Studies. *Energies* **2025**, *18*, 4452. [CrossRef]

14. Lucentini, I.; Colli, G.G.; Luzi, C.D.; Serrano, I.; Martínez, O.M.; Llorca, J. Catalytic ammonia decomposition over Ni-Ru supported on CeO2 for hydrogen production: Effect of metal loading and kinetic analysis. *Appl. Catal. B Environ.* **2021**, *286*, 119896. [CrossRef]
15. Meng, X.; Liu, M.; Mei, J.; Li, X.; Grigoriev, S.; Hasanien, H.M.; Tang, X.; Li, R.; Sun, C. Polarization loss decomposition-based online health state estimation for proton exchange membrane fuel cells. *Int. J. Hydrogen Energy* **2025**, *157*, 150162. [CrossRef]
16. Wu, X.; Cao, B.; Li, X.; Xu, J.; Ren, X. Component sizing optimization of plug-in hybrid electric vehicles. *Appl. Energy* **2011**, *88*, 799–804. [CrossRef]
17. Ezzat, M.F.; Dincer, I. Development and assessment of a new hybrid vehicle with ammonia and hydrogen. *Appl. Energy* **2018**, *219*, 226–239. [CrossRef]
18. Isikveren, A.T.; Kaiser, S.; Pornet, C.; Vratny, P.C. A Conceptual Design and Sizing Methodology for Hybrid-Electric Aircraft. *J. Aircr.* **2014**, *51*, 1793–1808.
19. Zhou, J.; Jiang, J.; Fan, F.; Sun, C.; Dong, Z.; Song, K. Real-Time Impedance Detection for PEM Fuel Cell Based on TAB Converter Voltage Perturbation. *Energies* **2024**, *17*, 4320. [CrossRef]
20. Martins, J.R.R.A.; Lambe, A.B. Multidisciplinary Design Optimization: A Survey of Architectures. *AIAA J.* **2013**, *51*, 2049–2075. [CrossRef]
21. Lu, D.; Yi, F.; Hu, D.; Li, J.; Yang, Q.; Wang, J. Online optimization of energy management strategy for FCV control parameters considering dual power source lifespan decay synergy. *Appl. Energy* **2023**, *348*, 121516. [CrossRef]
22. Simpson, T.W.; Peplinski, J.; Koch, P.N.; Allen, J.K. Kriging Models for Global Approximation in Simulation-Based Design. *AIAA J.* **2001**, *39*, 2233–2241. [CrossRef]
23. Li, Y.; Lu, X.; Kar, N.C. Rule-Based Control Strategy With Novel Parameters Optimization Using NSGA-II for Power-Split PHEV Operation Cost Minimization. *IEEE Trans. Veh. Technol.* **2014**, *63*, 3051–3061. [CrossRef]
24. Li, K.; Zhou, J.; Jia, C.; Yi, F.; Zhang, C. Energy sources durability energy management for fuel cell hybrid electric bus based on deep reinforcement learning considering future terrain information. *Int. J. Hydrogen Energy* **2024**, *52*, 821–833. [CrossRef]
25. Wen, Y.; Chu, Z.; Srivastava, A.; Teng, F.; Chen, B. Resilient Microgrid Scheduling With Synthetic Inertia From Electric Vehicles Within a Network of Charging Stations. *IEEE Trans. Control Syst. Technol.* **2024**, *33*, 775–787. [CrossRef]
26. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [CrossRef]
27. Pardhi, S.; El Baghdadi, M.; Hulsebos, O.; Hegazy, O. Optimal Powertrain Sizing of Series Hybrid Coach Running on Diesel and HVO for Lifetime Carbon Footprint and Total Cost Minimisation. *Energies* **2022**, *15*, 6974. [CrossRef]
28. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. [CrossRef]
29. Zhou, Q.; Zhang, W.; Cash, S.; Olatunbosun, O.; Xu, H.; Lu, G. Intelligent sizing of a series hybrid electric power-train system based on Chaos-enhanced accelerated particle swarm optimization. *Appl. Energy* **2017**, *189*, 588–601. [CrossRef]
30. Chen, B.; Li, X.; Evangelou, S.; Lot, R. Joint Propulsion and Cooling Energy Management of Hybrid Electric Vehicles by Optimal Control. *IEEE Trans. Veh. Technol.* **2020**, *69*, 4894–4906. [CrossRef]
31. Baydin, A.G.; Pearlmutter, B.A.; Radul, A.A.; Siskind, J.M. Automatic Differentiation in Machine Learning: A Survey. *J. Mach. Learn. Res.* **2018**, *18*, 1–43.
32. Tang, X.; Shi, L.; Li, M.; Xu, S.; Sun, C. Health State Estimation and Long-Term Durability Prediction for Vehicular PEM Fuel Cell Stacks Under Dynamic Operational Conditions. *IEEE Trans. Power Electron.* **2025**, *40*, 4498–4509. [CrossRef]
33. Ruan, S.; Ma, Y.; Wei, Z.; Zhang, C.; Xiang, C. Data-Driven Cooperative Differential Game Based Energy Management Strategy for Hybrid Electric Propulsion System of a Flying Car. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 14770–14784. [CrossRef]
34. Li, J.; Wang, W.; Chen, W.; Gong, Q.; Luo, J.; Lin, R.; Xin, H.; Zhang, H.; Wang, D.; Peng, Q.; et al. Sub-nm ruthenium cluster as an efficient and robust catalyst for decomposition and synthesis of ammonia: Break the "size shackles". *Nano Res.* **2018**, *11*, 4774–4785. [CrossRef]
35. Xie, P.; Yao, Y.; Huang, Z.; Liu, Z.; Zhang, J.; Li, T.; Wang, G.; Shahbazian-Yassar, R.; Hu, L.; Wang, C. Highly efficient decomposition of ammonia using high-entropy alloy catalysts. *Nat. Commun.* **2019**, *10*, 4011. [CrossRef]
36. Xu, K.; Jiang, N.; Wang, P.; Wang, W.W.; Jia, C.J. Highly efficient cobalt catalysts promoted by CeO_2-Al_2O_3 for ammonia decomposition. *Catal. Sci. Technol.* **2024**, *14*, 5678–5686. [CrossRef]
37. Makepeace, J.W.; Wood, T.J.; Hunter, H.M.A.; Jones, M.O.; David, W.I.F. Ammonia decomposition catalysis using non-stoichiometric lithium imide. *Chem. Sci.* **2015**, *6*, 3805–3815. [CrossRef] [PubMed]
38. Zhang, H.; Lei, N.; Liu, S.; Fan, Q.; Wang, Z. Data-driven predictive energy consumption minimization strategy for connected plug-in hybrid electric vehicles. *Energy* **2023**, *283*, 128514. [CrossRef]
39. Lei, N.; Zhang, H.; Hu, J.; Hu, Z.; Wang, Z. Sim-to-real design and development of reinforcement learning-based energy management strategies for fuel cell electric vehicles. *Appl. Energy* **2025**, *393*, 126030. [CrossRef]

40. Song, K.; Hou, T.; Jiang, J.; Grigoriev, S.A.; Fan, F.; Qin, J.; Wang, Z.; Sun, C. Thermal management of liquid-cooled proton exchange membrane fuel cell: A review. *J. Power Sources* **2025**, *648*, 237227. [CrossRef]
41. Deb, K.; Jain, H. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Trans. Evol. Comput.* **2014**, *18*, 577–601. [CrossRef]
42. Lei, N.; Zhang, H.; Wang, H.; Wang, Z. An Improved Co-Optimization of Component Sizing and Energy Management for Hybrid Powertrains Interacting With High-fidelity Model. *IEEE Trans. Veh. Technol.* **2023**, *72*, 15585–15596. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.