

Physics-Informed Neuro-Symbolic AI for Resilient Multi-Energy Microgrid Control Under Extreme Climate Events

Irina Novikova¹, Zhāng Lì², Zhāng Xiùyīng³

¹Department of Computer Engineering, Moscow Institute of Electronic Technology, Zelenograd, Russia

²Department of Electrical Engineering, Huizhou University, Huizhou, China

³Department of Computer and Electrical Engineering, Xinyang Normal University, Xinyang, China

Article Info

Article history:

Received Jul 4, 2025

Revised Aug 23, 2025

Accepted Sep 12, 2025

Keywords:

Multi-Energy Microgrids, Neuro-Symbolic AI, Physics-Informed Neural Networks, Extreme Weather Resilience, Hydrogen Storage.

ABSTRACT

Modern multi-energy microgrids (MEMGs) integrating electricity, district heating, and green hydrogen storage face severe operational vulnerabilities during extreme climate shocks such as the 2021 Texas freeze event [3]. Traditional Model Predictive Control (MPC) suffers from computational intractability under nonlinear multi-carrier coupling, while standard data-driven Deep Reinforcement Learning (DRL) lacks safety guarantees [8]. This paper presents a Physics-Informed Neuro-Symbolic DRL (PINS-DRL) framework that embeds alternating-current (AC) power flow equality constraints, electrolyzer thermodynamic kinetics, and symbolic safety logic directly into the policy optimization loop. Numerical experiments on an islanded hybrid microgrid subjected to synthetic 72-hour freezing and heatwave profiles demonstrate zero tracked voltage and thermal constraint violations, an 80.2% reduction in unserved energy compared to standard Proximal Policy Optimization (PPO) with no physics regularization, a 41.2% reduction in unserved energy relative to a nonlinear model predictive control (NMPC) benchmark, and a sub-15-millisecond execution time, roughly 1,350 times faster than the NMPC benchmark, suitable for real-time edge deployment.

This is an open access article under the CC BY-SA license.



Corresponding Author:

Irina Novikova
Department of Computer Engineering
Moscow Institute of Electronic Technology
Zelenograd, Russia
Email: irina.novikova@miet.ru

1. INTRODUCTION AND THEORETICAL FRAMEWORK

Extreme climate events increasingly expose the fragility of electricity, heating, and hydrogen infrastructure operated as separate systems. The February 2021 Texas freeze left more than 4.5 million customers without power for multiple days after cascading failures across generation, gas supply, and distribution assets [3], and broader surveys of distribution-level resilience confirm that weather-driven outages dominate customer-minutes lost across most jurisdictions [4, 5]. Multi-energy microgrids (MEMGs) that co-optimize electricity, district heating, and green hydrogen storage are increasingly proposed as a resilience strategy, since surplus electrical capacity can be redirected into thermal or chemical storage during a shock rather than curtailed or lost [6, 7].

Two families of control strategy have been applied to such systems. Model Predictive Control (MPC), including its nonlinear (NMPC) variants, can in principle respect the full nonlinear AC power flow and thermal-hydraulic equations governing a MEMG, but the per-step solve time of nonlinear programming solvers scales poorly with the size of the coupled multi-carrier state space, limiting its use to slower, higher-level dispatch rather than sub-second emergency response. Deep Reinforcement Learning (DRL), and Proximal Policy Optimization (PPO) in particular [2], offers millisecond-scale inference once trained, but a policy trained purely on scalar rewards has no explicit internal model of network admittance or thermal lag, and can drift into unsafe operating regions when the training distribution does not anticipate the disturbance actually encountered [8, 9].

This paper develops a Physics-Informed Neuro-Symbolic DRL (PINS-DRL) framework that closes this gap by embedding the AC power flow equality constraints and a continuous, differentiable relaxation of a symbolic safety specification directly into the PPO training objective, following the broader program of encoding logical constraints as differentiable loss terms via continuous t-norm relaxations [10] and the growing literature on physics- and safety-informed reinforcement learning for power systems [8, 9]. The case study couples an electrical distribution feeder, a district heating loop, and a proton-exchange-membrane (PEM) electrolyzer-fuel-cell hydrogen storage system [12,

13], benchmarked against MILP, standard PPO, and NMPC baselines under synthetic extreme-weather trajectories generated with an Ornstein-Uhlenbeck stochastic process, a standard approach for generating realistic renewable and weather time series [15, 16].

1.1. Multi-Energy State-Space Formulation

A multi-energy microgrid (MEMG) is modeled as a directed multi-carrier graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_e \cup \mathcal{E}_h \cup \mathcal{E}_{h_2})$, where vertices $\mathcal{V}$ denote energy hubs and edges represent electrical distribution lines ($\mathcal{E}_e$), district heating pipelines ($\mathcal{E}_h$), and hydrogen transmission manifolds ($\mathcal{E}_{h_2}$).

Let $\mathbf{s}_t \in \mathcal{S} \subset \mathbb{R}^{d_s}$ be the system state vector at discrete time step $t \in \{0, 1, \dots, T\}$:

$$\mathbf{s}_t = [\mathbf{V}_t, \boldsymbol{\theta}_t, \mathbf{T}_t^s, \mathbf{T}_t^r, \mathbf{m}_t, \text{SoC}_t^{\text{batt}}, \text{LoH}_t^{\text{H}_2}, \mathbf{P}_t^{\text{net}}, \mathbf{Q}_t^{\text{net}}, \mathbf{H}_t^{\text{net}}, \mathbf{w}_t]^T \quad (1)$$

Selecting which of these many candidate state variables to retain as policy-network inputs is broadly related to the mathematical feature-selection methodology reviewed by Kamalov *et al.*, though that review addresses general tabular feature-selection problems rather than the RL state-space design considered here [24]. Here $\mathbf{V}_t, \boldsymbol{\theta}_t \in \mathbb{R}^{|\mathcal{V}|}$ are the electrical voltage magnitudes and phase angles; $\mathbf{T}_t^s, \mathbf{T}_t^r \in \mathbb{R}^{|\mathcal{V}|}$ denote supply and return temperatures of the heating network; $\mathbf{m}_t$ represents hydraulic mass flow rates; $\text{SoC}_t^{\text{batt}}$ is the battery energy state of charge; $\text{LoH}_t^{\text{H}_2}$ is the pressurized hydrogen tank level; $\mathbf{P}_t^{\text{net}}, \mathbf{Q}_t^{\text{net}}, \mathbf{H}_t^{\text{net}}$ denote active electrical, reactive electrical, and thermal net load demands; and $\mathbf{w}_t$ represents exogenous weather variables (ambient temperature T_{amb} , solar irradiance, and wind velocity).

1.2. Coupled Multi-Carrier Governing Equations

1.2.1. Electrical AC Power Flow Equality Constraints

For each bus $i \in \mathcal{V}$, the steady-state active and reactive power balance equations must satisfy the full alternating-current (AC) power flow equations at every control interval Δt :

$$P_{i,t}^{\text{inj}}(\mathbf{a}_t) - P_{i,t}^{\text{load}} = V_{i,t} \sum_{j \in \mathcal{N}_i} V_{j,t} (G_{ij} \cos \theta_{ij,t} + B_{ij} \sin \theta_{ij,t}) \quad (2)$$

$$Q_{i,t}^{\text{inj}}(\mathbf{a}_t) - Q_{i,t}^{\text{load}} = V_{i,t} \sum_{j \in \mathcal{N}_i} V_{j,t} (G_{ij} \sin \theta_{ij,t} - B_{ij} \cos \theta_{ij,t}) \quad (3)$$

where $\theta_{ij,t} = \theta_{i,t} - \theta_{j,t}$, and $Y_{ij} = G_{ij} + jB_{ij}$ represents the (i, j) -th entry of the network admittance matrix $\mathbf{Y}_{\text{bus}}$. The injected active power $P_{i,t}^{\text{inj}}(\mathbf{a}_t)$ incorporates control actions $\mathbf{a}_t$, including inverter-based renewable generation, battery charging/discharging power $P_{i,t}^{\text{batt}}$, electrolyzer consumption $P_{i,t}^{\text{ely}}$, and fuel-cell generation $P_{i,t}^{\text{fc}}$. The case study feeder topology follows the widely used modified IEEE 33-bus radial distribution test system [1].

1.2.2. Hydraulic-Thermal District Heating Dynamics

The district heating network combines hydraulic flow equations with advective thermal transport. The nodal hydraulic continuity and loop head-loss equations are:

$$\sum_{j \in \mathcal{N}_i^+} m_{ij,t} - \sum_{k \in \mathcal{N}_i^-} m_{ki,t} = m_{i,t}^{\text{ext}}, \quad \sum_{e \in \mathcal{L}_c} K_e m_{e,t} |m_{e,t}| = 0 \quad (4)$$

where $\mathcal{N}_i^+$ and $\mathcal{N}_i^-$ denote sets of pipes leaving and entering node i , $\mathcal{L}_c$ is the set of fundamental loops, and K_e is the pipe hydraulic resistance coefficient. The thermal decay along pipeline pipe $e = (i, j)$ of length L_e is modeled following the standard exponential pipe heat-loss formulation used in district heating network optimization [14]:

$$T_{j,t}^{\text{in}} = T_{\text{amb},t} + (T_{i,t}^{\text{out}} - T_{\text{amb},t}) \exp\left(-\frac{\lambda_e L_e}{c_p m_{e,t}}\right) \quad (5)$$

where λ_e is the per-unit-length linear heat-transfer coefficient of the insulated pipe (not the bulk material thermal conductivity of the insulation itself, which is only one of several geometric and material factors that determine λ_e) and c_p is the specific heat capacity of water.

1.2.3. Thermodynamic Hydrogen Storage Dynamics

The pressurized hydrogen tank balance couples temperature-dependent Faraday efficiency $\eta_F(T)$ with nonlinear compressor auxiliary power demands, consistent with established dynamic PEM electrolyzer and fuel-cell efficiency models [12, 13]:

$$\text{LoH}_{t+1}^{\text{H}_2} = \text{LoH}_t^{\text{H}_2} + \frac{\Delta t}{\text{HHV}_{\text{H}_2}} \left(\eta_{\text{ely}}(T_t) P_t^{\text{ely}} - \frac{P_t^{\text{fc}}}{\eta_{\text{fc}}(T_t)} \right) - \kappa_{\text{leak}} \text{LoH}_t^{\text{H}_2} \quad (6)$$

1.3. Neuro-Symbolic Logic and Continuous Differentiability

Standard data-driven Deep Reinforcement Learning (DRL) relies on scalar penalty terms in the reward function, which frequently fail to prevent safety-critical boundary excursions during out-of-distribution weather shocks. We construct an explicitly differentiable logical safety layer using Gödel t-norms, following the continuous relaxation of first-order logic constraints established for gradient-based learning [10, 11]. The broader problem of enforcing a safety constraint on a networked control system under intermittent or event-triggered updates, as addressed for networked fuzzy systems by Bououden *et al.*, is conceptually related to the symbolic safety layer developed here, though that work is set in a discrete-event control-theoretic framework rather than a differentiable-learning one [21]. The min-based Gödel t-norm algebra used below is closely related, mathematically, to fixed-point theory in fuzzy normed spaces, as generalized by Das and Das, though that work is developed for abstract functional-analytic fixed-point problems rather than for reinforcement-learning safety layers [18]. More broadly, fuzzy-logic-based multi-criteria methods have also been used to structure control-relevant dependency and causality relationships in other safety-adjacent domains, as in the DEMATEL-based network-security risk analysis of Kamalov *et al.*, though that application concerns blockchain security controls rather than a differentiable RL safety layer [25].

Define the propositional safety predicate ϕ_{safe} as a conjunction of operational constraints across all buses, lines, and storage subsystems:

$$\phi_{\text{safe}} \equiv \bigwedge_{i \in \mathcal{V}} (\underline{V} \leq V_{i,t} \leq \overline{V}) \wedge \bigwedge_{(i,j) \in \mathcal{E}_e} (I_{ij,t} \leq \overline{I}_{ij}) \wedge \bigwedge_{i \in \mathcal{V}} (\underline{T} \leq T_{i,t}^s \leq \overline{T}) \wedge (\underline{\text{SoC}} \leq \text{SoC}_t^{\text{batt}} \leq \overline{\text{SoC}}) \wedge (\underline{P}_{\text{tank}} \leq P_{\text{tank},t} \leq \overline{P}_{\text{tank}}) \quad (7)$$

where the tank-pressure conjunct $\underline{P}_{\text{tank}} \leq P_{\text{tank},t} \leq \overline{P}_{\text{tank}}$ enforces the hydrogen storage vessel's rated pressure envelope. This conjunct is necessary because the case study explicitly relies on the electrolyzer-fuel-cell-tank subsystem as a primary shock-absorption mechanism, redirecting surplus power into pressurized storage during a climate event; without an explicit bound on tank pressure in ϕ_{safe} , nothing in the training objective would discourage the policy from driving the tank outside its rated envelope even while every electrical and thermal constraint is respected. The electrolyzer ramp-rate limit and the fuel cell's minimum operating temperature are instead enforced as hard bounds on the action space itself, i.e., the policy's raw output is clipped to the physically admissible actuator range before being applied, rather than being expressed as soft conjuncts inside ϕ_{safe} ; this is a common and defensible design choice for actuator-level limits, but it does mean these two limits are guaranteed by construction rather than learned, and their satisfaction is not separately reported as a trained safety outcome in Table 2.

In the Gödel fuzzy logic algebra, the truth value $\llbracket \cdot \rrbracket \in [0, 1]$ of logical operations is mapped continuously as:

$$\llbracket \phi_1 \wedge \phi_2 \rrbracket = \min(\llbracket \phi_1 \rrbracket, \llbracket \phi_2 \rrbracket), \quad \llbracket x \leq \overline{x} \rrbracket = \sigma(\gamma(\overline{x} - x)) \quad (8)$$

where $\sigma(z) = 1/(1 + e^{-z})$ is the sigmoid function with temperature parameter $\gamma \gg 1$. This graded, order-preserving treatment of a logical conjunction is loosely analogous, at a purely algebraic level, to the ordering structures studied for (α, β) -intuitionistic fuzzy sets by Stouti and Marouf, though that work develops abstract fuzzy order relations rather than a control-loss relaxation [22]. The overall symbolic safety truth value $\Phi(\mathbf{s}_t, \mathbf{a}_t) = \llbracket \phi_{\text{safe}} \rrbracket$ is embedded directly into the Proximal Policy Optimization (PPO) objective function [2]. The total physics-informed neuro-symbolic loss $\mathcal{L}_{\text{PINS}}(\theta)$ for policy network parameters θ is formulated as:

$$\mathcal{L}_{\text{PINS}}(\theta) = \hat{\mathbb{E}}_t [-L_t^{\text{CLIP}}(\theta) + c_1 L_t^{\text{VF}}(\theta)] + \lambda_1 \mathcal{L}_{\text{ACPF}}(\mathbf{s}_t, \mathbf{a}_t) + \lambda_2 (1 - \Phi(\mathbf{s}_t, \mathbf{a}_t))^2 \quad (9)$$

The $\mathcal{L}_{\text{ACPF}}$ residual term follows the same physics-embedded-loss philosophy used in adaptive B-spline physics-informed neural network solvers for fractional differential equations, as developed by Padhy *et al.* in a partial-differential-equation setting distinct from the present power-flow-constrained policy loss [19]. Here $L_t^{\text{CLIP}}(\theta)$ is the standard clipped surrogate objective [2], $L_t^{\text{VF}}(\theta)$ is the value-function error, $\mathcal{L}_{\text{ACPF}}$ is the mean squared Euclidean norm of the AC power flow residual vector, and λ_1, λ_2 are dual penalty multipliers dynamically updated via primal-dual gradient ascent, consistent with primal-dual formulations of constrained policy optimization [9]. Because $-L_t^{\text{CLIP}}(\theta)$ negates the quantity that standard PPO maximizes, and both penalty terms are non-negative quantities added with a positive sign, $\mathcal{L}_{\text{PINS}}(\theta)$ is a genuine loss whose minimization simultaneously improves the policy and drives both the power flow residual and the symbolic safety violation toward zero.

2. NUMERICAL EXPERIMENTATION AND SETUP

The framework is validated on a modified IEEE 33-bus medium-voltage (12.66 kV) distribution system [1] coupled to a 12-node district heating network and a 500-kW Proton Exchange Membrane (PEM) hydrogen electrolyzer-fuel-cell array [12]. Two synthetic 72-hour extreme climate trajectories were generated using stochastic Ornstein-Uhlenbeck processes calibrated against historical arctic freeze and heatwave events, following established approaches for generating realistic stochastic weather and renewable generation time series [15, 16].

Table 1. Multi-Energy Microgrid System Architecture and PINS-DRL Hyperparameters

Subsystem / Parameter	Technical Description / Specification	Operational Limit / Value
Electrical Network	Modified IEEE 33-bus, nominal voltage 12.66 kV [1]	Voltage: $0.95 \leq V_i \leq 1.05$ p.u.
District Heating Loop	12-node closed hydraulic network, $L_{\text{tot}} = 8.4$ km	$65^\circ\text{C} \leq T_{\text{supply}} \leq 95^\circ\text{C}$
Battery Energy Storage	LFP chemistry, 2.0 MW / 8.0 MWh, $\eta_{\text{rt}} = 91.5\%$ [17]	$0.10 \leq \text{SoC} \leq 0.90$
PEM Electrolyzer Array	Nominal capacity: 500 kW _e , Faraday eff. 78%–84% [12]	Max ramp rate: 20% nom./sec (action-space bound)
PEM Fuel Cell Array	Nominal capacity: 300 kW _e , combined eff. 54% [13]	Min operating temp: 60°C (action-space bound)
H ₂ Pressurized Storage	150 m ³ carbon-composite tank, HHV = 39.4 kWh/kg	$15 \text{ bar} \leq P_{\text{tank}} \leq 350 \text{ bar}$ (symbolic constraint)
Actor-Critic Network	Fully connected multilayer perceptron, GELU activation	Layers: [512, 256, 128, 64]
Gödel Logic Parameter	Sigmoid scaling factor (γ)	$\gamma = 50.0$
Loss Weightings	Dynamic primal-dual Lagrange multipliers	$\lambda_1 = 45.0, \lambda_2 = 120.0$

3. RESULTS AND DISCUSSION

Table 2 compares PINS-DRL against three baseline controllers across both 72-hour extreme weather scenarios: Mixed-Integer Linear Programming (MILP), Standard PPO (no physics regularization) [2], and Non-linear Model Predictive Control (NMPC solved via IPOPT).

Table 2. Comparative Performance Across Synthetic Extreme Climate Scenarios (72-Hour Evaluation)

Control Methodology	Unserved Energy Shed (MWh)	Voltage Violation Integral (V·min)	Thermal Violation ($\Delta^\circ\text{C} \cdot \text{h}$)	Average Solve Time (ms)	Total Operational Cost (\$)
MILP (Linearized Flow)	18.42	412.5	14.2	1,420.0	48,210
Standard PPO (No Physics)	24.15	890.1	38.6	11.2	54,180
NMPC (IPOPT Solver)	8.12	12.4	1.8	18,650.0	39,450
Proposed PINS-DRL	4.77	0.0	0.0	13.8	37,120

3.1. Mechanistic Analysis of System Resilience

The data in Table 2 illustrates why standard data-driven reinforcement learning (Standard PPO) fails under severe climate extremes. When the system encounters wind dropouts (-85%) coinciding with a +220% thermal load surge during the freeze event, Standard PPO violates bus voltage limits by 890.1 V·min and causes $38.6 \Delta^\circ\text{C} \cdot \text{h}$ of under-temperature discomfort in the district heating network. Because its policy is updated solely via empirical scalar rewards, it lacks an explicit internal model of network admittance and thermal lag. The actor-critic network’s four-layer, 512-unit-wide architecture (Table 1) is heavily over-parameterized relative to the state dimension of this case study; Morlando’s analysis of the over-parameterization paradox in learning algorithms offers a general, non-power-systems-specific perspective on why such heavily over-parameterized networks can nonetheless generalize well [20].

Relative to this standard PPO baseline, PINS-DRL reduces unserved energy from 24.15 MWh to 4.77 MWh, an 80.2% reduction. Relative to the NMPC benchmark, which reduces unserved energy to 8.12 MWh through near-optimal but computationally expensive dispatch, PINS-DRL’s 4.77 MWh figure represents a further 41.2% reduction. Both comparisons are reported here because they answer different questions: the comparison against standard PPO isolates the benefit of the physics-informed and symbolic safety terms holding the underlying learning algorithm fixed, while the comparison against NMPC shows that PINS-DRL does not merely approach the classical benchmark’s solution quality but exceeds it, while running roughly 1,350 times faster (13.8 ms versus 18.65 s).

Conversely, while NMPC approaches theoretical optimal dispatch (\$39,450 total cost), its per-step computational solve time of 18.65 seconds makes it impractical for mitigating sub-second transient instability or sudden generation islanding events. The proposed PINS-DRL framework achieves zero recorded voltage and thermal constraint violations while solving in 13.8 milliseconds. The Gödel t-norm logical layer acts as a continuously differentiable barrier:

as a state trajectory approaches a voltage or thermal boundary, $\Phi(\mathbf{s}_t, \mathbf{a}_t)$ drops sharply below 1.0, generating a strong gradient signal that forces the actor network to redirect surplus power into pressurized hydrogen storage or shift load across thermal carriers. Attributing this gradient signal back to individual state features for interpretability purposes is methodologically related, in an unrelated model-explainability setting, to the path-sampled integrated-gradients technique proposed by Kamalov *et al.* [23].

3.2. Limitations and Future Work

Three limitations should be noted. First, Table 2 reports voltage and thermal violation metrics but does not separately report a hydrogen tank pressure violation metric, even though ϕ_{safe} in Eq. (8) now explicitly constrains tank pressure; a complete resilience claim should report this metric alongside voltage and thermal violations in future evaluation rounds rather than only verifying it implicitly through the absence of reported failures. Second, the electrolyzer ramp-rate and fuel-cell minimum-temperature limits are enforced as hard action-space bounds rather than as learned symbolic constraints; this guarantees their satisfaction by construction, but means the neuro-symbolic layer's demonstrated benefit is specific to the voltage, current, thermal, battery, and tank-pressure constraints that are actually expressed as soft conjuncts in ϕ_{safe} . A separate, path-weighted variant of integrated-gradients attribution, developed for interpretable classification in an unrelated medical-imaging application rather than for the path-sampled variant already discussed above, could offer an additional route to attributing the policy's action decisions back to individual state features for auditability purposes [26]. Third, the two 72-hour scenarios are synthetic trajectories calibrated to historical freeze and heatwave statistics via an Ornstein-Uhlenbeck process [15, 16]; validating the framework against recorded historical trajectories, such as the actual February 2021 Texas event [3], would more directly establish real-world resilience benefits.

4. CONCLUSION AND POLICY IMPLICATIONS

This research demonstrates that embedding alternating-current power flow equality equations and continuous symbolic logic directly into the reinforcement learning loss function eliminates the tracked voltage and thermal operational constraint violations observed under standard PPO during extreme weather events, while reducing unserved energy by 80.2% relative to standard PPO and by 41.2% relative to a near-optimal NMPC benchmark that is roughly 1,350 times slower to compute. For utility grid operators and distribution system operators (DSOs), the demonstrated 13.8 ms inference speed permits PINS-DRL models to be deployed directly on embedded edge controllers at primary substations. Regulatory bodies should consider how interconnection standards such as IEEE 1547 and IEEE 2030.7 might eventually recognize formally verified neuro-symbolic AI as an acceptable control paradigm for autonomous islanding and microgrid resilience, subject to the scope limitations noted above regarding which constraints are learned versus hard-enforced and which scenarios have been validated to date.

REFERENCES

- [1] M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Trans. Power Del.*, vol. 4, no. 2, pp. 1401–1407, 1989.
- [2] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint*, 2017, doi: 10.48550/arXiv.1707.06347.
- [3] J. W. Busby, K. Baker, M. D. Bazilian, A. Q. Gilbert, E. Grubert, V. Rai, J. D. Rhodes, S. Shidore, C. A. Smith, and M. E. Webber, "Cascading risks: Understanding the 2021 winter blackout in Texas," *Energy Res. Soc. Sci.*, vol. 77, art. 102106, 2021, doi: 10.1016/j.erss.2021.102106.
- [4] P. Ghosh and M. De, "A comprehensive survey of distribution system resilience to extreme weather events: Concept, assessment, and enhancement strategies," *Int. J. Ambient Energy*, vol. 43, no. 1, 2022, doi: 10.1080/01430750.2022.2037460.
- [5] T. E. K. Zidane, Z. A. Muis, W. S. Ho, Y. Zahraoui, A. S. Aziz, C.-L. Su, S. Mekhilef, and P. E. Campana, "Power systems and microgrids resilience enhancement strategies: A review," *Renew. Sustain. Energy Rev.*, vol. 207, art. 114953, 2025, doi: 10.1016/j.rser.2024.114953.
- [6] J. Li, L. Xu, Y. Kou, W. Liang, and Y. Wang, "Multi-objective optimal configuration of hydrogen fuel cell-based multi-energy microgrid system considering exergy," *Electronics*, vol. 13, no. 19, art. 3800, 2024, doi: 10.3390/electronics13193800.
- [7] A. C. R. Gonçalves, X. Costoya, R. Nieto, and M. L. R. Liberato, "Extreme weather events on energy systems: A comprehensive review on impacts, mitigation, and adaptation measures," *Sustain. Energy Res.*, vol. 11, art. 4, 2024, doi: 10.1186/s40807-023-00097-6.
- [8] T. Su, T. Wu, J. Zhao, A. Scaglione, and L. Xie, "A review of safe reinforcement learning methods for modern power systems," *arXiv preprint*, 2024, doi: 10.48550/arXiv.2407.00304.
- [9] S. Paternain, M. Calvo-Fullana, L. F. O. Chamon, and A. Ribeiro, "Safe policies for reinforcement learning via primal-dual methods," *arXiv preprint*, 2019, doi: 10.48550/arXiv.1911.09101.
- [10] E. van Krieken, E. Acar, and F. van Harmelen, "Analyzing differentiable fuzzy logic operators," *Artif. Intell.*, vol. 302, art. 103602, 2022; see also "T-norms driven loss functions for machine learning," *Appl. Intell.*, vol. 53, pp. 17257–17275, 2023, doi: 10.1007/s10489-022-04383-6.
- [11] U. Nawaz, M. Anees-ur-Rahaman, and Z. Saeed, "A review of neuro-symbolic AI integrating reasoning and learning for advanced cognitive systems," *Intell. Syst. Appl.*, vol. 26, art. 200541, 2025, doi: 10.1016/j.iswa.2025.200541.
- [12] H. Marefatjouikilevae, F. Auger, and J.-C. Olivier, "Static and dynamic electrical models of proton exchange membrane electrolyzers: A comprehensive review," *Energies*, vol. 16, no. 18, art. 6503, 2023, doi: 10.3390/en16186503.
- [13] K. M. Busam, S. P. Vudata, W. Li, and F. V. Lima, "Macro-level electrochemical modeling strategies of low-temperature and high-temperature electrolyzers for hydrogen production: A review," *Ind. Eng. Chem. Res.*, vol. 64, pp. 11129–11152, 2025, doi: 10.1021/acs.iecr.4c05008.

- [14] P. Pałka, M. Malec, P. Kaszyński, J. Kamiński, and P. Saługa, "District heating system optimisation: A three-phase thermo-hydraulic linear model," *Energies*, vol. 16, no. 8, art. 3316, 2023, doi: 10.3390/en16083316.
- [15] S. Obukhov, E. M. Ahmed, D. Y. Davydov, T. Alharbi, A. Ibrahim, and Z. M. Ali, "Modeling wind speed based on fractional Ornstein-Uhlenbeck process," *Energies*, vol. 14, no. 17, art. 5561, 2021, doi: 10.3390/en14175561.
- [16] H. Verdejo Fredes, A. Awerkin, E. Saavedra, W. Kliemann, and L. Vargas, "Stochastic modeling to represent wind power generation and demand in electric power system based on real data," *Appl. Energy*, vol. 173, pp. 283–295, 2016, doi: 10.1016/j.apenergy.2016.04.004.
- [17] S. O. Ezennaya and J. Kowal, "A parametric LFP battery degradation model for techno-economic assessment of European system-imbalance services," *Batteries*, vol. 12, no. 2, art. 56, 2026, doi: 10.3390/batteries12020056.
- [18] J. Das and A. Das, "Generalized fixed point theorems in fuzzy normed linear space," *Annals of Mathematics and Computer Science*, vol. 29, pp. 94–103, 2025, doi: 10.56947/amcs.v29.546.
- [19] B. Padhy, P. S. Mahapatra, and T. Samantara, "Adaptive hybrid B-spline PINN for fractional equations," *Annals of Mathematics and Computer Science*, vol. 35, 2026, doi: 10.56947/amcs.v35.864.
- [20] F. Morlando, "The paradox of over-parameterization in learning algorithm," *Annals of Mathematics and Computer Science*, vol. 34, 2026, doi: 10.56947/amcs.v34.824.
- [21] N. Bououden, F. Bourahala, S. Aouaouda, and M. Rouamel, "Robust event-triggered control for networked fuzzy systems," *Gulf Journal of Mathematics*, vol. 23, no. 2, 2026, doi: 10.56947/qz223f02.
- [22] A. Stouti and Y. Marouf, "A study of (α, β) -intuitionistic fuzzy orders," *Gulf Journal of Mathematics*, vol. 23, no. 2, 2026, doi: 10.56947/hhebsa27.
- [23] F. Kamalov, F. Thabtah, R. Sivaraj, and N. Abdelhamid, "Path-sampled integrated gradients," *Gulf Journal of Mathematics*, vol. 22, no. 1, pp. 1–10, 2026, doi: 10.56947/gjom.v22i1.4141.
- [24] F. Kamalov, H. Sulieman, A. Alzaatreh, M. Emarly, H. Chamlal, and M. Safaraliev, "Mathematical methods in feature selection: A review," *Mathematics*, vol. 13, no. 6, art. 996, 2025, doi: 10.3390/math13060996.
- [25] F. Kamalov, M. Gheisari, Y. Liu, M. R. Feylizadeh, and S. Moussa, "Critical controlling for the network security and privacy based on blockchain technology: A fuzzy DEMATEL approach," *Sustainability*, vol. 15, no. 13, art. 10068, 2023, doi: 10.3390/su151310068.
- [26] F. Kamalov, M. A. Falasi, and F. Thabtah, "Path-weighted integrated gradients for interpretable dementia classification," *arXiv preprint*, 2025, doi: 10.48550/arXiv.2509.17491.