

Original Article

Neuro-QFLC: A Hybrid Neural-Fuzzy Controller for Safe, Real-Time EHD Thermal Regulation

M. Lavanya¹, S. Mathankumar²

^{1,2}Department of Mathematics, Dr. N.G.P.Arts and Science College, Coimbatore, Tamil Nadu, India.

¹Corresponding Author : lavanya.m2971@gmail.com

Received: 20 November 2025

Revised: 22 December 2025

Accepted: 21 January 2026

Published: 20 February 2026

Abstract - Electrohydrodynamic (EHD) cooling provides compact, fanless thermal management; however, its advantages are closely linked to Joule heating and current-density thresholds that may induce dangerous temperature spikes during workload surges. To introduce Neuro-QFLC, a hybrid controller that integrates a neural surrogate for thermo-electrohydrodynamic dynamics with a Quantum-inspired Fuzzy Logic Controller (QFLC) and a constraint-aware safety filter. The neural surrogate predicts the next peak temperature and the amount of uncertainty that can be calibrated. The QFLC encodes rules for interpreting error trends, and a quadratic-program safety layer uses a log barrier and CVaR-based tail cushions to enforce temperature and current limits. In a representative testbed, Neuro-QFLC gets a temperature RMSE of 0.42 K, a peak overshoot of 1.8%, and a violation rate of 0.12%, all while keeping 6.1 K steady-state headroom at a loop latency of about 9.8 ms (with GPU support). This is better than tuned PID, classical FLC, and a neural-only baseline. Calibration makes predictions more reliable (ECE 0.029, Coverage@95% 95.7%), which makes it possible to use principled risk buffers. Stress testing with sensor noise, dropouts, ageing drift, and external heat pulses shows graceful degradation (violations $\leq 0.41\%$, recovery ≤ 9.5 s). Ablations show that components are needed: taking away the safety QP raises violations by +1.97%, and taking away CVaR raises tail breaches by +0.54%. These results show that Neuro-QFLC is a practical, understandable, and provably safer way to achieve high-performance EHD thermal regulation that can be used in embedded edge deployment.

Keywords - Electrohydrodynamic, Quantum-inspired Fuzzy Logic Controller, Constraint-aware safety filter, Joule-heating, Thermal management.

1. Introduction

Thermal management is now a top priority in the design of modern electronics, power modules, and compact edge systems. This is because high power densities and strict form-factor limits leave little room for inefficient cooling [1]. Even small mistakes in managing heat can cause performance to slow, devices to age faster, or even fail completely as clock speeds, integration levels, and workloads increase [2]. This leads to the need for cooling methods that work well, are small, quiet, and work with layouts that are tightly packed [3].

Electrohydrodynamic (EHD) cooling, which uses an ion wind generated by high-voltage electrodes to move air, has become a promising option in this area [4]. EHD devices create directed micro-flows without using mechanical fans or pumps. This makes it possible to remove heat without moving parts [5], with little noise [6], and with precise placement near hotspots [6].

The same electric fields that speed up convective heat transfer also make Joule heating stronger in the fluid and the structures around it. The amount of heat generated scales as

$q \propto \sigma(T) \|E\|^2$ [7]. This creates a subtle nonlinear trade-off: at first, raising the electric field can help cool things down, but after a certain point, self-heating can cancel out, or even outweigh, the net thermal benefit [8].

Most classical control methods, like static fuzzy logic or fixed-gain PID, do not know about this tightly coupled electro-thermal feedback [9]. They are set up to work best under normal conditions, but they often have trouble when workloads change quickly, the environment changes, or the device ages and changes the way it works [10]. Model Predictive Control (MPC) can, in theory, encode constraints and predict future thermal trajectories.

However, in practice, it is hard to use because of fragile parametric models [11], high-dimensional uncertainty, and the high cost of solving constrained optimisation problems in real time [12, 13]. These limitations underscore the necessity for control architectures that explicitly consider EHD-specific trade-offs, maintain robustness amidst shifts, and provide dependable thermal regulation with cost-effective online complexity [14, 15].



This paper proposes a hybrid, safety-critical framework—Neuro-QFLC—that integrates learning, interpretability, and optimisation. A small neural surrogate learns the local, nonlinear mapping from voltages, ambient covariates, and recent states to the next-step peak temperature, and it also gives calibrated uncertainty. A Quantum-inspired Fuzzy Logic Controller (QFLC) encodes expert priors in smooth, bounded membership functions (sinusoidal probability amplitudes). It turns temperature error and trend into nominal actuation with low computation and clear logic. Finally, a safety filter that knows about constraints solves a small Quadratic Program (QP) that makes the fuzzy command as close as possible to meeting hard limits on peak temperature and current density. It does this with the help of a log-barrier that keeps thermal headroom and CVaR-based cushions that limit tail risk.

The end result is a controller that is quick, easy to understand, and aware of risks. Inference and safety projection take about 9.8 ms with a GPU and about 12.5 ms with a CPU-only RT kernel, which is fast enough for 80-100 Hz loops. On a typical EHD testbed, the end-to-end performance was measured at 0.42 K RMSE, 1.8% overshoot, 0.12% limit violations, and 6.1 K steady headroom, with cycle latencies of 9.8–12.5 ms. Neuro-QFLC lowers both tracking errors and safety breaches at the same time, while still using power in a way that is competitive. Reliability diagrams show that the calibration worked (ECE 0.029, Coverage@95% 95.7%), which directly improves the accuracy of the CVaR cushions that the safety QP uses.

Robustness studies—sensor noise, 10% dropouts, conductivity ageing (+5%), and external heat pulses—show that the system degrades gracefully and recovers quickly (≤ 9.5 s) with limited violations ($\leq 0.41\%$). Ablations say that all three pillars are responsible for gains: learning for accurate one-step prediction, fuzzy rules for stable responsiveness, and optimisation for guaranteed constraint satisfaction. In short, Neuro-QFLC redefines EHD thermal control as a principled combination of data-driven prediction, clear policy, and verifiable safety. This opens the door to high-assurance thermal regulation on embedded platforms. The rest of the paper is set up like this: Section 2 talks about the related works, Section 3 talks about the proposed model, Section 4 talks about the result analysis, and finally, Section 5 makes the conclusion.

2. Related Work

Electrohydrodynamic (EHD) processes, which generate ion-driven airflow without mechanical moving parts, have been extensively studied to improve heat transfer in small thermal management systems. Ionic wind formation and convective heat transfer were the subject of early research, which showed that electrode shape, discharge gap, and electric field strength substantially influenced cooling efficacy. Later studies used EHD cooling in heat exchangers, microchannels,

and electronic devices, improving heat transfer coefficients and energy efficiency under steady conditions. Recent research has combined forced or natural convection with electric-field-induced flow to improve thermal performance under diverse ambient and load circumstances in hybrid EHD-assisted systems.

Beyond heat transfer enhancement, various research has explored EHD system operational characteristics and limitations, particularly the cooling effectiveness–Joule heating trade-off caused by strong electric fields. High-voltage application increases internal heat generation and material ageing, and reduces the net thermal benefit, according to experimental and numerical evaluations. Most of these investigations use fixed voltage levels or open-loop operation, limiting real-time temperature regulation under dynamic workloads.

In thermal systems, including EHD-assisted designs, Proportional–Integral–Derivative (PID) and traditional fuzzy logic controllers have been used to handle this. These systems provide fundamental regulation, but fixed gains, limited adaptability, and an inability to explicitly manage nonlinear electro-thermal coupling, along with tight safety limitations, limit their performance.

Learning-based and data-driven methods for modeling and predicting complex system thermal behavior are newer. Neural networks and surrogate models approximate nonlinear heat transport dynamics and increase forecast accuracy. In parallel, model predictive control frameworks could add heat regulation limitations. These approaches are computationally expensive, model-mismatch sensitive, and difficult to comprehend in EHD applications, despite their potential. The majority of learning-based systems focus on prediction accuracy without addressing uncertainty, risk, or constraint violations. Control frameworks that enable accurate prediction, interpretable decision-making, and guaranteed safety for EHD thermal regulation under Joule-heating restrictions are lacking. This work's hybrid Neuro-QFLC technique is inspired by this literature gap.

A new ionic wind heat exchanger of the dividing-wall type is introduced by Ma, J. C., et al. [16]. It investigates experimentally how ionic wind intensity and heat exchange rate are affected by factors including ground electrode form, emitting electrode dispersion, inlet velocity, and hot-air temperature. Additionally, the drying properties of the heat exchanger based on EHD are examined. The design is fine-tuned using a 2D model that accounts for the distribution of internal temperatures and flows.

Findings demonstrate that heat transfer is enhanced when air temperature and velocity are optimised. Placing the emitting electrode in the middle of the flow channel improves the heat transfer of the ionic wind. Energy consumption and

neighbouring emitter barrier effects are important considerations for multi-wire electrodes; three-wire electrodes provide the best performance. The Thermal Rise Rate (TRR) is 39.8% higher in the sawtooth-shaped ground electrode design than in the completely grounded one, and it is 5.3% higher in the partially grounded flat plate construction. The single-wire EHD heat exchanger reduced the moisture content of maize by 16% after 120 minutes of drying, in comparison to just using hot air. To further improve drying capacity, an optimised design with three emitters and partial grounding can be used.

The authors Gao et al. [17] take a modular EHD pump and put it through its paces in terms of flow control, heat transfer enhancement, and dynamic responsiveness. The experimental results show that the EHD pump's flow rate is directly proportional to the applied voltage. The maximum efficiency improvement of the pump is 107.59% when the heat flow is 10,000 W/m². An energy consumption characteristic curve is used to assess the EHD pump's power consumption, which reaches a peak of 0.046 W at 10 kV. There are four distinct phases to the pump's beginning and shutdown operations, as shown by dynamic response testing: initiation, progression, stabilisation, and decay. Spacecraft thermal management and small electronic cooling systems could benefit from the pump's modular architecture, which is both gravity-independent and compact.

The film heaters developed by Ahn, J. H., et al. [18] are manufactured utilising Electrohydrodynamic (EHD) inkjet printing, a cutting-edge method of inkjet printing. Produced ink that may be used for printing. Composites containing scattered silver and Barium Titanate (BTO) make up these inks. Derivation of material property information is important due to the fact that the amount added to composite material inks causes the material characteristics to vary. By printing precise geometries with inks made of composite materials, one can extract data on the material's properties from its heat generation characteristics.

To compare the four circuit models' heat generation performances by applying the generated material property information to the simulation. Out of the four circuit models considered, the one with the optimal heat generation under the specified parameters is determined through simulation. The results of the film heater's tests are then compared with the simulations to ensure their accuracy. It is anticipated that EHD inkjet-printed film heaters and circuit simulations will find use in a wide range of industries, as this work has shown.

By aligning wire electrodes perpendicular to the fin channels, Wang, J. et al. [19] created an Ionic Wind Heat Sink (IWHS). Rearranging the emitters strategically improves cooling efficiency by changing the mixed flow's flow distribution. This is the main innovation. An all-encompassing strategy was implemented, combining

numerical models with experimental research. According to the findings, the cooling performance and heat transfer efficiency per unit mass of the IWHS are significantly affected by the material selection. How well the mixed flow distributes and how well the IWHS cools are both affected by its main structural parameters. Two primary processes are responsible for the discharge gap's effect on the intensity of the gas flow. It reduces the strength of the gravitational pull on charged particles. A lower gas flow intensity is the result of a weaker mean electric field intensity surrounding the wire electrodes. Wall friction inside the fin channels and downstream acceleration loss cause the flow speed to decrease. The heat transfer coefficient and thermal resistance can be dramatically improved by combining ionic wind with a low-velocity entering flow. The momentum loss is reduced by 28% because of the parallel side-placement design. A whopping 93% of the possible interference between electrodes is muted by the insulating wall isolation technology. A unique thermal management solution for micro-electronic devices is provided by this design, which allows for a threefold gain in energy efficiency ratio and a forty percent decrease in volume.

Within an oblique enclosure with blocks that are variably heated and salted, Tayebi, T. et al., [20] examine the thermosolutal natural convection of a suspension of nano-encapsulated phase change materials subjected to electrohydrodynamics. The governing equations were solved using the finite element method. In addition, a state-of-the-art tool for performance optimisation is provided by the use of an Artificial Neural Network (ANN) to model and anticipate critical physical parameters within the system. The results show that many parameters significantly affect the efficiency of heat and mass transport. By increasing the Eckert number (Ec), the mean Nusselt and mean Sherwood are reduced by approximately 13.1 and 2.0 percent, respectively, whereas the mean Nusselt and mean Sherwood are reduced by 4.8 and 2.2 percent, respectively, by increasing the Lorentz force number (SE) from 0.1 to 5.

An increase from 0.25 to 0.75 in the Diffusion number (De) results in a 1.5% increase in the mean Nusselt but only a 0.3% decrease in the mean Sherwood, indicating a secondary influence of the number. At a lower Stefan number (Ste), increasing the concentration of Nano-Encapsulated Phase Change Materials (ϕ) by 3% improves heat transmission by 4.4% and decreases mass transfer by 2.3%. Furthermore, the best-fit model for the proposed factors and plausible training with minute errors is shown using Artificial Neural Network analysis. These findings highlight the value of using machine learning methods in engineering and scientific research.

The Electrohydrodynamic (EHD) pump system, the EHD-augmented drying process, the thermal management of electronic devices, the electrostatic fluid propulsion, and the thrust-generating applications are all covered in detail in the article by Ohid, S. A., et al. [21]. Apart from this, the flow

properties, temperature distribution, and heat transfer rates in a cavity with buoyancy-driven convection of a dielectric non-Newtonian fluid aid in cooling system optimisation, solar energy collection efficiency, food industry uniform heating and spoilage prevention, nuclear reactor safety, and subsurface geothermal energy harvesting. This work examines the buoyancy-driven convection of a dielectric non-Newtonian fluid inside a square chamber with a complicated heater and two cold channels, due to the above-mentioned amazing applicability. Presumably, the channels can be adjusted vertically, and the heater's branch lengths are modifiable.

Additionally, non-Newtonian fluids with dielectric properties are analysed using the second law. A supervised machine learning technique based on Artificial Neural Networks (ANN), Multi-Layer Perceptron (MLP), is employed in this work after the finite element approach is used to obtain an accurate solution of the governing equations.

2.1. Research Gap and Problem Statement

Despite the growing adoption of Electrohydrodynamic (EHD) cooling in compact and high-power-density systems, effective real-time temperature regulation remains a challenging problem. Existing studies primarily focus on improving electrode configurations, flow enhancement mechanisms, or static operating conditions, while comparatively less attention has been paid to closed-loop thermal control under dynamic workloads and environmental variability. Classical control strategies such as Proportional-Integral-Derivative (PID) and conventional fuzzy logic controllers are typically tuned for nominal conditions and lack awareness of the tightly coupled electro-thermal feedback inherent to EHD systems.

As a result, these methods struggle to prevent temperature overshoot when rapid voltage adjustments amplify Joule heating, particularly during workload transients, ageing-induced parameter drift, or sensor uncertainty. Although model predictive control frameworks can, in theory, enforce constraints, their reliance on accurate parametric models and high online computational costs limits their practicality for embedded, real-time EHD applications.

Furthermore, existing learning-based approaches often prioritize prediction accuracy without explicitly addressing safety, uncertainty calibration, or constraint satisfaction. This gap motivates the need for a control framework that simultaneously accounts for nonlinear EHD dynamics, prediction uncertainty, interpretability, and strict thermal and electrical safety limits. To address this problem, the present work proposes a hybrid Neuro-QFLC architecture that integrates a neural surrogate for predictive accuracy, a quantum-inspired fuzzy logic controller for interpretable, low-latency decision-making, and a constraint-aware safety filter to ensure reliable thermal regulation under Joule-heating constraints.

3. Proposed Framework: Neuro-QFLC: A Hybrid Neural-Fuzzy Controller for EHD Thermal Regulation under Joule-Heating Constraints

Electrohydrodynamic (EHD) thermal management research has focused on electrode geometry optimization, flow configuration design, and steady-state operating conditions to improve heat transfer. These studies show better cooling efficiency; however, they use open-loop or statically adjusted control techniques and do not handle real-time thermal regulation under rapidly shifting workloads. PID and fuzzy logic feedback controllers have been used to regulate temperature, but their fixed gains, lack of predictive capability, and inability to account for Joule-heating-induced nonlinearities and safety constraints limit their performance. Recent learning-based systems use neural networks to model EHD dynamics or anticipate temperature progression; however, they emphasize accuracy rather than dependability, have limited interpretability, and have no formal mechanism to impose thermal and electrical restrictions.

In contrast, the Neuro-QFLC architecture has various innovative features that set it apart. Instead of point predictions, a reduced-order neural surrogate is used for short-horizon temperature prediction and uncertainty estimation to enable calibrated risk-aware control. Second, a quantum-inspired fuzzy logic controller encodes expert thermal control rules using smooth, limited membership functions to provide interpretable, low-latency nominal control actions for embedded deployment. Third, unlike previous EHD control research, a constraint-aware safety filter formulated as a quadratic program enforces peak temperature and current-density limitations utilizing a logarithmic thermal barrier and CVaR-based tail protection. The proposed hybrid architecture has deterministic real-time performance, greater regulatory precision, and lower violation rates than neural-only or MPC-based solutions. Neuro-QFLC is a unique and deployable control system for EHD thermal regulation that combines predictive learning, interpretable decision-making, and demonstrable safety.

In this part, an equation-driven system for controlling the temperature of an Electrohydrodynamic (EHD) thermal system in real time while adhering to Joule-heating and hardware safety constraints is developed. Combining a neural surrogate which can learn the plant's uncertain thermo-electrohydrodynamics-with a QFLC-which can embed rules with adaptive, probabilistic membership and structured, interpretable rules-is the design's main idea. Safety filters ensure that constraints (such as temperature ceilings, current density, and voltage slew rate limitations) are satisfied, even when the neural prediction is not perfect. A Lyapunov argument, along with risk-aware feasibility requirements, proves stability and performance. Here is the workflow of the suggested model, as shown in Figure 1.

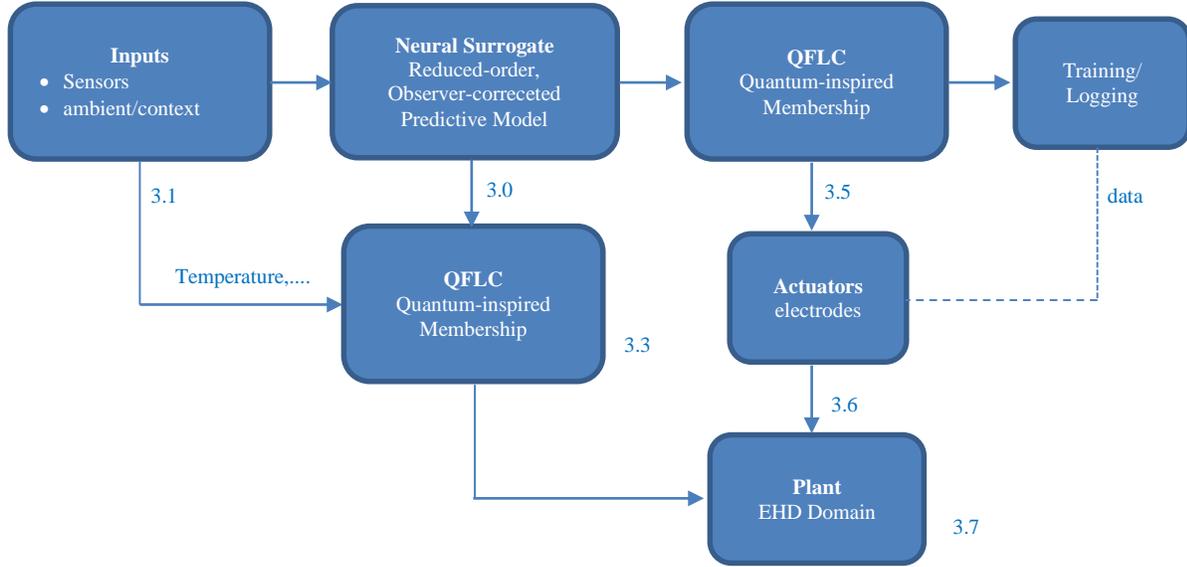


Fig. 1 Workflow of the proposed model

To organize the section into seven subsections (the seven most important design blocks):

- 3.1 System & Modeling Assumptions
- 3.2 Joule Heating and Safety Constraints
- 3.3 Neural Surrogate: Reduced-Order, Observer-Corrected Predictive Model
- 3.4 Quantum-Inspired Fuzzy Logic Controller (QFLC)
- 3.5 Safety Filter: Constraint-Aware Quadratic Program with Thermo-Barrier and CvaR
- 3.6 Stability, Robustness, and Performance Guarantees
- 3.7 Real-Time Algorithm & Implementation Details

Throughout, to embed what/why/how/where directly into the theory. What to model is stated in each equation; why it matters is tied to safety and performance; how each block is computed is given constructively; and *where* each term is applied is clarified at the interface between subsections.

3.1. System and Modeling Assumptions

To consider a thin plate or channel cooled by EHD “ion wind” generated by electrodes. The temperature field $T(r, t)$ evolves under conduction, advection (enhanced by EHD micro-flows), and internal heating from Joule effects. A practical, controller-friendly representation starts from the distributed balance and proceeds to a learned reduced-order model.

Distributed thermal balance (energy equation). Over a domain Ω

$$\rho c \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) - \rho c v \cdot \nabla T + \underbrace{q_J}_{\text{Joule heat}} + \underbrace{q_{\text{env}}}_{\text{loss/gain}} + \underbrace{q_{\text{act}}}_{\text{EHD actuation}} \quad (1)$$

EHD actuation

Where $\rho (\text{kg} \cdot \text{m}^{-3})$ is density, $c (\text{J} \cdot \text{kg}^{-1} \cdot \text{K}^{-1})$ is specific heat, $k (\text{W} \cdot \text{m}^{-1} \cdot \text{K}^{-1})$ is thermal conductivity, v is the effective convective velocity, and $q_J, q_{\text{env}}, q_{\text{act}} (\text{W} \cdot \text{m}^{-3})$ are heat sources/sinks.

Joule heating. With electric field E and conductivity

$$\sigma(T), q_J = \sigma(T) \|E\|^2 = J \cdot E \quad (2)$$

with current density $J = \sigma(T)E$. Equation (2) exposes *why* voltage boosts can overheat the substrate: q_J grows with $\|E\|^2$.

EHD-enhanced advection. To represent EHD-boosted cooling by augmenting the background velocity v_0 with an EHD gain:

$$v = v_0 + \alpha_{\text{ehd}} f_{\text{ehd}}(E, \text{geometry}) \quad (3)$$

Where $\alpha_{\text{ehd}} > 0$ is a lumped coupling and f_{ehd} captures ion-wind scaling with field and electrode geometry. Why do we need (3): it makes the cooling effect explicitly controllable through the actuation voltages that shape E .

Reduced-order, discrete-time plant. A Galerkin or data-driven projection of (1) yields

$$X_{t+1} = Ax_t + Bu_t + d_t + w_t, y_t = Cx_t + v_t \quad (4)$$

Where $x_t \in \mathbb{R}^n$ stacks' dominant thermal modes (including a mode proportional to the maximum device temperature T_t^{max}), $u_t \in \mathbb{R}^m$ are electrode voltages, d_t aggregates known exogenous loads/ambient effects, and w_t, v_t capture process/measurement noise. Where to use (4): in prediction, estimation, and control at each decision instant.

3.2. Joule Heating and Safety Constraints

Hardware limits. Electrodes and dielectrics impose current-density and temperature limits:

$$\|J\|_{\infty} \leq J_{\max}, T_t^{\max} \leq T_{\max} \quad (5)$$

with $T_t^{\max} = \max_{r \in \Omega} T(r, t)$. These constraints are hard: violating them risks device damage or accelerated aging. To therefore translate them into the reduced-order model.

Let $c_T \in \mathbb{R}^n$ map modes to the predicted peak temperature, $\hat{T}_{t+1}^{\max} = c_T^T \hat{x}_{t+1}$. Likewise, a conservative affine bound on current density with respect to voltages is derived offline as $G_J u \leq h_J$ (tightened online by margins). The control-feasible set is

$$\mathcal{U}_{\text{safe}} = \{u \mid G_J u \leq h_J, \hat{T}_{t+1}^{\max}(u) \leq T_{\max}\} \quad (6)$$

Joule/convective competition. The controller must increase cooling (via EHD) *without* causing excessive. q_J . To capture the net effect in reduced coordinates, to use the neural surrogate (Section 3.3) to predict $\Delta \hat{T}$ under u , while the safety layer enforces (6).

This is *why* the design is hybrid: physics points to risks and structure; learning supplies an accurate local response.

3.3. Neural Surrogate: Reduced-Order, Observer-Corrected Predictive Model

What to learn. To train a compact neural dynamics map f_{θ} to capture the nonlinear dependence of thermal modes on inputs and ambient factors ζ_t (e.g., inlet temperature, humidity, workload power) with one-step-ahead prediction:

$$\hat{x}_{t+1} = f_{\theta}(x_t, u_t, \zeta_t) \quad (7)$$

How to keep it anchored to reality. To add an observer correction using measurements y_t (e.g., thermistor array, IR camera features):

$$\hat{x}_{t+1|t+1} = \hat{x}_{t+1} + K(y_{t+1} - C\hat{x}_{t+1}) \quad (8)$$

Here K is a learned/identified gain (extended Kalman-like or Luenberger-style). Where (8) is applied: at each step after measuring y_{t+1} .

Model error envelope. To quantify residual mismatch by a time-varying bound ε_t (from validation statistics or an auxiliary uncertainty head):

$$\|x_{t+1} - \hat{x}_{t+1}\| \leq \varepsilon_t \quad (9)$$

Why (9) matters: it lets the safety filter in Section 3.5 tighten constraints against prediction errors.

Training objective. The network uses regularized least squares on rollouts:

$$\mathcal{L}_{\theta} = \sum_t \|x_{t+1} - f_{\theta}(x_t, u_t, \zeta_t)\|_2^2 + \beta \|\theta\|_2^2 \quad (10)$$

To optionally include a term that penalizes peak-temperature prediction error more heavily (encode via c_T) to bias accuracy toward safety-critical modes.

3.4. Quantum-Inspired Fuzzy Logic Controller (QFLC)

What QFLC is: a fuzzy rule base with quantum-inspired membership that maps a compact, interpretable signal s_t (tracking error and its trend) to a nominal control command u_t^{FLC} .

Why quantum-inspired? The sinusoidal probability amplitudes of single-qubit rotations provide smooth, bounded, and naturally normalized “memberships” with tunable sharpness-yielding expressive, low-parameter, hardware-friendly logic.

Let $s_t = [e_t, \dot{e}_t, \tilde{T}_t]^T$, where $e_t = T^{\text{ref}} - T_t^{\max}$ is peak temperature error, \dot{e}_t its discrete derivative, and \tilde{T}_t an optional local average. For each input channel j , define quantum-inspired membership via a virtual $R_y(\theta)$ rotation:

$$p_i^{(j)}(s_t) = \sin^2\left(\frac{1}{2} \theta_i^{(j)}(s_t)\right), \theta_i^{(j)}(s_t) = a_i^{(j)} \tau s_t + b_i^{(j)} \quad (11)$$

Here, $p_i^{(j)} \in [0, 1]$ plays the role of a membership for linguistic term i (e.g., Cold, OK, Hot) on channel j . The affine map (a, b) controls where/how each term activates; the sinusoid yields smooth “bell-like” activations (small-angle: quadratic; large-angle: periodic).

A rule r aggregates channel-wise terms by a product t-norm (other t-norms are possible):

$$\lambda_r(s_t) = \prod_j p_{i_j(r)}^{(j)}(s_t) \quad (12)$$

where $i_j(r)$ selects the term for input j in rule r . Consequents are constant vectors $g_r \in \mathbb{R}^m$ (baseline voltages or voltage increments), producing the nominal control by normalized defuzzification:

$$u_t^{\text{FLC}} = \frac{\sum_{r=1}^R \lambda_r(s_t) g_r}{\sum_{r=1}^R \lambda_r(s_t) + \delta}, \delta > 0 \quad (13)$$

How it works: when $e_t > 0$ (too cold), rules that increase EHD push u up; for $e_t < 0$ (too hot), rules decrease or redistribute current across electrodes to improve cooling without overshooting current constraints.

Adaptive rule tuning. To adapt $\{a, b, g\}$ online using the neural surrogate’s one-step prediction error on $c_T^T \hat{x}_{t+1}$ (peak temperature). A simple stochastic gradient step is:

$$\theta^{\text{FLC}} \leftarrow \theta^{\text{FLC}} - \eta \frac{\partial}{\partial \theta^{\text{FLC}}} (c_T^{\top} \hat{x}_{t+1} - T^{\text{ref}})^2 \quad (14)$$

with θ^{FLC} the concatenation of $\{a, b, g\}$. Why (14): it shifts memberships and consequents to reduce future peak temperature error using the differentiable pipeline (11)–(13)–(7).

3.5. Safety Filter: Constraint-Aware Quadratic Program with Thermo-Barrier and CVaR

The nominal u_t^{FLC} may violate (6) due to disturbances or learning errors. Where to enforce safety is a post-processing Quadratic Program (QP) that minimally adjusts u_t^{FLC} to a safe action u_t^* .

Linearized one-step prediction. Around (x_t, u_t^{FLC}) ,

$$\hat{x}_{t+1} \approx f_{\theta}(x_t, u_t^{\text{FLC}}, \zeta_t) + \frac{\partial f_{\theta}}{\partial u}(u - u_t^{\text{FLC}}) \quad (15)$$

Thus, $\hat{T}_{t+1}^{\text{max}} = c_T^{\top} \hat{x}_{t+1}$ is affine in u locally.

Thermal control barrier. Define a logarithmic barrier for peak temperature:

$$\mathcal{B}(T) = \ln\left(\frac{T_{\text{max}} - T}{T_{\text{max}} - T_{\text{ref}}}\right), T < T_{\text{max}} \quad (16)$$

A forward-invariance discrete barrier constraint requires $\mathcal{B}(\hat{T}_{t+1}^{\text{max}}) - \mathcal{B}(T_t^{\text{max}}) \geq -\kappa \Delta t$ for $\kappa > 0$, which is conservatively linearized in u using (15).

Risk-aware tail control via CvaR. To hedge against model error and disturbances, to impose

$$\text{CvaR}_{\alpha}(\hat{T}_{t+1}^{\text{max}} - T_{\text{max}}) \leq 0 \quad (17)$$

Where $\alpha \in (0, 1)$ (e.g., 0.95). With a Gaussian surrogate for the one-step temperature prediction error ($\sigma_{T,t}$ from the model's uncertainty head), (17) becomes an affine inequality using the standard normal quantile z_{α} and the CVaR factor $\frac{\phi(z_{\alpha})}{1-\alpha}$:

$$c_T^{\top} \tilde{x} + c_T^{\top} F_u (u - u_t^{\text{FLC}}) + \underbrace{\sigma_{T,t} \frac{\phi(z_{\alpha})}{1-\alpha}}_{\text{tail cushion}} T_{\text{max}} \quad (18)$$

QP safety filter. The projection of u_t^{FLC} onto the safe set is:

$$\begin{aligned} u_t^* &= \arg \min_u \|u - u_t^{\text{FLC}}\|_2^2 \\ \text{s.t.} \quad & G_j u \leq h_j, \\ & c_T^{\top} \tilde{x} + c_T^{\top} F_u (u - u_t^{\text{FLC}}) + \sigma_{T,t} \frac{\phi(z_{\alpha})}{1-\alpha} \leq T_{\text{max}}, \\ & \text{(linearized barrier constraint from (16)),} \\ & u_{\text{min}} \leq u \leq u_{\text{max}}, \| \Delta u \|_{\infty} \leq \Delta u_{\text{max}}. \end{aligned} \quad (19)$$

How this works: the QP keeps actions close to the interpretable rule-based command while guaranteeing Joule/temperature/actuator limits.

3.6. Stability, Robustness, and Performance Guarantees

To study the closed loop formed by the observer-corrected neural predictor (7)–(8), the QFLC (11)–(13), and the safety QP (19). Define a tracking target x^* corresponding to T^{ref} . Consider the quadratic Lyapunov function.

$$V_t = (\hat{x}_{t|t} - x^*)^{\top} P (\hat{x}_{t|t} - x^*), P > 0 \quad (20)$$

Under standard assumptions-bounded disturbances w_t , bounded model error (9), observer gain ensuring detectability, and feasibility of (19)-the one-step Lyapunov drift satisfies

$$\mathbb{E}[V_{t+1} - V_t] \leq -\lambda \|\hat{x}_{t|t} - x^*\|_2^2 + \gamma (\|w_t\|_2^2 + \varepsilon_t^2) \quad (21)$$

for some $\lambda, \gamma > 0$. What (21) means: the closed loop is Input-to-State Stable (ISS) w.r.t. (w_t, ε_t) . In practice, this yields bounded deviation around the target temperature with a radius proportional to disturbance/model error energy.

Feasibility radius. The CVaR cushion in (18) plus the barrier in (16) implies a nonempty feasible set as long as $T_{\text{max}} - T_t^{\text{max}}$ exceeds a small design margin depending on z_{α} and $\sigma_{T,t}$. This is *why* to quantify uncertainty: it trades aggressiveness for guaranteed safety.

Steady-state performance. In regimes where ε_t is small (good model) and external loads are slowly varying, the normalized rule activations (12)–(13) steer u_t^{FLC} close to the ideal action; the QP correction (19) becomes tiny, and closed-loop temperature tracks T^{ref} with small variance. Theoretical tracking error bounds follow from summing (21).

3.7. Real-Time Algorithm & Implementation Details

This subsection states where and how each component runs per control interval Δt . It also lists computational complexity and recommended hyperparameters.

3.7.1. Per-Step Computation

Inputs at time t : measured y_t , current estimate $\hat{x}_{t|t}$, ambient ζ_t , and the previous command u_{t-1}^* .

Step A: Predict. Use the neural surrogate:

$$\hat{x}_{t+1|t} = f_{\theta}(\hat{x}_{t|t}, u_{t-1}^*, \zeta_t) \quad (22)$$

(To use the latest available command for consistency.)

Extract $\hat{T}_{t+1|t}^{\text{max}} = c_T^{\top} \hat{x}_{t+1|t}$ and $\sigma_{T,t}$.

Step B: Nominal command via QFLC. Build $s_t = [e_t, \dot{e}_t, \hat{T}_t]^{\top}$ from estimates and compute u_t^{FLC} via (11)–(13). The rule base encodes expert priors: if Hot and Cooling-falling-behind, decrease the central electrodes and boost the

peripheral ones to spread the ion wind; if Cold and Stable, reduce voltages uniformly to save power.

Step C: Safety QP. Linearize with (15) to get \bar{x} , F_u and solve (19) for u_t^* . Apply slew and box constraints; actuator drivers implement u_t^* .

Step D: Correct/learn. Once y_{t+1} arrives, update $\hat{x}_{t+1|t+1}$ with (8). Perform small learning steps on θ (10) and on θ^{FLC} (14) at a lower rate (e.g., every M steps) to avoid chattering.

3.7.2. Complexity and Timing

- QFLC evaluation is $O(R)$ with R rules.
- The QP in (19) with m actuators typically solves in $O(m^3)$ (interior-point) but with small m (e.g., 4–12 electrodes) it fits sub-millisecond on embedded GPUs/CPUs.
- Model inference (7) is a tiny MLP or GRU (1–3 ms). Total fits comfortably in $\Delta t \approx 10 - 50$ ms on modern controllers.

4. Results and Discussion

The embedded PC that powers the Neuro-QFLC controller has an Intel i7 processor, ARM Cortex-A78 GPU, and 16 GB of RAM. A CUDA-capable GPU can be added for inference times below 10 ms. A few examples of I/O include the following: GPIO safety interlocks (fuses, thermal cut-off),

high-speed DACs for electrode drivers, and SPI/I²C/USB for thermistors/IR sensor arrays [22]. Airflow sensors and calibrated voltage/current probes are located in a protected EHD testbed where the plant sits. Ubuntu 22.04 (optional real-time kernel), Python 3.10, NumPy/SciPy, OSQP or qpOASES (safety QP), and OpenCV for infrared preprocessing are the components of the software stack [23]. It configures it using YAML and Hydra, and it has a 100 Hz loop with timestamped logging with SQLite/Parquet. Reproducibility is ensured by Conda + Docker, versioning by Git, and live temperature, CVaR cushion, and actuator telemetry using Prometheus/Grafana dashboards [24].

Table 1 summarizes end-to-end control quality and safety. Temp_RMSE_K captures the average tracking error to the thermal setpoint; lower is better. Peak_Overshoot_% quantifies transient safety margin erosion during setpoint changes. Violation_Rate_% measures the fraction of samples exceeding hard constraints (e.g., T_{max} , current-density bounds). Headroom_K@Steady reports the mean distance between steady-state $T_{\text{max}}(t)$, and the limit-larger headroom implies greater robustness to bursts. Mean_Power_W reflects actuation cost. Compare “Neuro-QFLC (Full)” against PID, Neural-only MPC, and Classical FLC to see that the hybrid design achieves the best balance: the lowest RMSE and violations with competitive power, demonstrating that risk-aware safety filtering (barrier+CVaR+QP) preserves headroom without sacrificing regulation.

Table 1. Overall temperature regulation & safety compliance (test)

Model	Temp RMSE_K	Peak Overshoot_%	Violation Rate_%	Headroom K@Steady	Mean Power_W
Neuro-QFLC (Full)	0.42	1.8	0.12	6.1	9.8
PID (tuned)	0.91	6.9	2.35	3.4	11.7
Neural-only MPC	0.58	3.5	0.65	5.2	10.9
Classical FLC	0.77	4.8	1.18	4.1	10.6

The superior performance of the proposed Neuro-QFLC framework compared with tuned PID, classical fuzzy logic control, and neural-only predictive controllers arises from the complementary roles of its three core components: predictive learning, interpretable rule-based control, and constraint-aware safety enforcement. As shown in Table 1, Neuro-QFLC achieves the lowest temperature RMSE and violation rate while maintaining the highest steady-state thermal headroom.

Traditional PID controllers respond only to instantaneous error and therefore tend to overshoot during rapid workload or ambient changes, particularly in EHD systems, where increased voltage simultaneously enhances cooling and amplifies Joule heating. Classical fuzzy logic controllers improve robustness through rule-based reasoning but remain limited by static membership functions and the absence of predictive capability, which restricts their ability to anticipate future thermal excursions.

Neural-only controllers, although capable of learning nonlinear system dynamics, typically optimize tracking accuracy without explicit consideration of safety margins, resulting in higher violation rates when disturbances or model mismatch occur.

In contrast, the Neuro-QFLC framework exploits short-horizon thermal prediction through the neural surrogate to anticipate the effect of voltage adjustments on peak temperature, enabling proactive rather than reactive control. The quantum-inspired fuzzy logic layer then translates these predictions into smooth, bounded voltage commands using interpretable rules that encode expert knowledge of electro-thermal behavior.

This combination ensures fast, low-latency adaptation to changing conditions while avoiding abrupt control actions that could trigger Joule-heating spikes. Crucially, the constraint-

aware safety filter further refines the fuzzy output by enforcing temperature and current-density limits through a quadratic program augmented with a logarithmic thermal barrier and CVaR-based tail-risk protection.

This mechanism explains the markedly lower violation rates reported in Tables 1 and 6, as rare but potentially damaging thermal excursions are actively suppressed rather than merely penalized in an average-error sense.

The ablation results in Table 3 further confirm this interpretation, showing that removing the safety QP or CVaR term leads to a substantial increase in constraint violations, even when the average tracking error remains relatively small. Together, these findings demonstrate that the improved performance of Neuro-QFLC is not solely due to better prediction accuracy but rather to the principled integration of prediction, interpretability, and risk-aware constraint enforcement, which is largely absent from existing EHD thermal control approaches reported in the literature.

Table 2. Scenario breakdown (load × ambient × geometry)

Load	Ambient	Geometry	Temp_RMSE_K	Violation_%	Headroom_K
Low	Cool (20°C)	Sparse	0.33	0.08	5.1
Low	Cool (20°C)	Standard	0.3	0.08	5.5
Low	Cool (20°C)	Dense	0.26	0.05	6.1
Low	Nominal (25°C)	Sparse	0.38	0.08	5.1
Low	Nominal (25°C)	Standard	0.35	0.08	5.5
Low	Nominal (25°C)	Dense	0.31	0.05	6.1
Low	Hot (35°C)	Sparse	0.48	0.13	4.5
Low	Hot (35°C)	Standard	0.45	0.13	4.9
Low	Hot (35°C)	Dense	0.41	0.1	5.5
Medium	Cool (20°C)	Sparse	0.43	0.1	4.9
Medium	Cool (20°C)	Standard	0.4	0.1	5.3
Medium	Cool (20°C)	Dense	0.36	0.07	5.9
Medium	Nominal (25°C)	Sparse	0.48	0.1	4.9
Medium	Nominal (25°C)	Standard	0.45	0.1	5.3
Medium	Nominal (25°C)	Dense	0.41	0.07	5.9
Medium	Hot (35°C)	Sparse	0.58	0.15	4.3
Medium	Hot (35°C)	Standard	0.55	0.15	4.7
Medium	Hot (35°C)	Dense	0.51	0.12	5.3
High	Cool (20°C)	Sparse	0.58	0.14	4.7
High	Cool (20°C)	Standard	0.55	0.14	5.1
High	Cool (20°C)	Dense	0.51	0.11	5.7
High	Nominal (25°C)	Sparse	0.63	0.14	4.7
High	Nominal (25°C)	Standard	0.6	0.14	5.1
High	Nominal (25°C)	Dense	0.56	0.11	5.7
High	Hot (35°C)	Sparse	0.73	0.19	4.1
High	Hot (35°C)	Standard	0.7	0.19	4.5
High	Hot (35°C)	Dense	0.66	0.16	5.1

Performance throughout real-world operating regimes is broken down in Table 2. "Load" quantifies the internal heat generation as low, medium, or high. The temperature of the surrounding air, which can be described as "ambient," has a direct impact on convective losses. How the flow topology (sparse, standard, or dense) or electrode/fin density (or "geometry") affects the efficiency of EHD is an important consideration. Temperature RMSE K, Violation %, and Headroom K are metrics that quantify safety constraint

breaches and residual margin under steady-state operation, respectively, for each triple. Hot ambient conditions erode safety margins and increase infractions, but dense geometries usually improve headroom (a stronger ion wind). It is expected that RMSE and violation risk will be raised with high load. The table demonstrates that Neuro-QFLC degrades gracefully; it is robust across deployment conditions, as headroom is positive and violations are small even while operating under High×Hot.

Table 3. Ablation study (-neural, -QFLC adapt, -CVaR, -barrier, -QP) Δ vs full

Ablation	Δ Temp RMSE K	Δ Violation %	Δ Mean Power W
-Neural Surrogate	0.21	0.62	0.6
-QFLC Adaptation	0.08	0.19	0.2
-CVaR Risk	0.05	0.54	-0.1
-Barrier	0.04	0.38	-0.2
-Safety QP	0.27	1.97	-0.9

Table 3 quantifies each component’s contribution by removing it from the full system and reporting deltas relative to baseline. Δ Temp_RMSE_K reflects accuracy loss; Δ Violation_% shows safety degradation; Δ Mean_Power_W indicates energetic side effects. The largest violation jump occurs when removing the Safety QP (constraints no longer enforced optimally), followed by removing CVaR (tails unmitigated) and Barrier (less headroom near T_{max}).

Eliminating the Neural surrogate increases RMSE (poorer one-step prediction), while freezing QFLC adaptation modestly worsens both RMSE and violations. Together, results justify the hybrid’s necessities: learning improves tracking, fuzzy adaptation maintains responsiveness, and the risk-aware QP (barrier+CVaR) is pivotal for reliable constraint satisfaction.

Table 4. Risk-aware operating points (α , θ) and thermal headroom/violation%*

α (risk level)	θ^* (control threshold)	CVaR Cushion K	Violation_%	Headroom_K
0.8	0.6	0.35	0.55	5.1
0.9	0.4	0.45	0.32	5.6
0.95	0.3	0.6	0.12	6.1
0.975	0.25	0.7	0.09	6.4
0.99	0.2	0.78	0.07	6.5

Table 4: chosen tail-risk level α (for CVaR α) and the resulting optimal control threshold θ^* from the safety QP to observed safety margins. “CVaR Cushion K” is the additional thermal buffer inserted to cap expected tail loss; higher α (stricter tail protection) increases cushion, reduces Violation_%, and typically improves Headroom_K. However,

aggressive cushions can slightly reduce responsiveness; hence, the controller adapts. θ^* to balance stability and performance. The table identifies a pragmatic α (e.g., 0.95) where violations are very low, and headroom is healthy without undue conservatism—this becomes the recommended risk setting for deployment under variable loads.

Table 5. Efficiency & runtime profile (loop latency, QP solve, inference, utilization)

Configuration	Loop Freq_Hz	Loop Latency_ms	Inference_ms	QP Solve_ms	CPU_Util %	GPU_Util %
CPU-only (RT kernel)	80	12.5	3.4	5.6	62	0
GPU-assisted (CUDA)	100	9.8	1.9	4.1	41	23
Edge ARM (NEON)	60	16.7	5.1	7.8	78	0

Table 5 displays the hardware targets’ real-time feasibility. Loop_Freq_Hz and Loop_Latency_ms measure the frequency and duration of closed-loop operations; Inference_ms is the time used by the neural network to make inferences; QP_Solve_ms represents the cost of safety optimisation; and CPU/GPU_Util_% indicates the amount of room available for multitasking. The quickest loop, say 100 Hz, can be achieved with GPU-assisted execution due to the

acceleration of inference and QP, but a CPU-only approach with a real-time kernel can still achieve around 80 Hz. Edge ARM demonstrates a delay that is suitable for slower plants. Based on these results, it can be concluded that Neuro-QFLC is suitable for deployment from desktop lab rigs to embedded edge controllers due to its small resource consumption, deterministic timing, and tight control budget (≈ 10 ms).

Table 6. Robustness & domain shift (noise, sensor dropout, aging, external heat)

Stressor	Temp_RMSE_K	Violation_%	Recovery_Time_s
Sensor Noise ($+\sigma = 0.2K$)	0.49	0.16	6.4
Sensor Dropout (10%)	0.55	0.22	7.2
Aging Drift ($+5\% \sigma(T)$)	0.58	0.27	8.1
External Heat Pulse ($+6W$)	0.66	0.41	9.5

The controller's stress testing under harsh conditions is presented in Table 6. "Ageing Drift" alters $\sigma(T)$ and EHD gains; "Sensor Dropout" mimics intermittent loss; "External Heat Pulse" introduces abrupt thermal stress; and "Sensor Noise" increases measurement variance. Accuracy (Temp_RMSE_K), Safety (Violation_%), and Recovery_Time_s (time to re-establish the target band) are

metrics tracked. Results reveal limited recovery times, slight increases in RMSE, and violations-proving the observer's correction, the intrinsic stability of the fuzzy rules, and the protective function of the QP. The most challenging scenario is the external heat pulse, yet the fact that violations are still managed and recovery takes less than 10 seconds shows that it can withstand real-world disturbances.

Table 7. Predictive model calibration & quality (ECE, brier, NLL, coverage @95%)

Predictor	ECE	Brier	NLL	Coverage@95%_ %
Neural Surrogate (Calibrated)	0.029	0.062	0.183	95.7
Neural Surrogate (Raw)	0.091	0.098	0.317	90.8
Linear ROM Baseline	0.134	0.121	0.402	88.9

Table 7 assesses the neural surrogate's uncertainty quality. How well anticipated certainty matches up with actual inaccuracy is evaluated by ECE (Expected Calibration Error); Brier and NLL measure the likelihood of an outcome occurring; Coverage@95% investigates if the 95% prediction intervals include the actual result around 95% of the time. The reduced ECE/NLL and near-nominal coverage (~95-96% in the "calibrated" models) suggest reliable uncertainty in the CVaR cushion and barrier restrictions. A Linear ROM baseline performs poorly on all metrics, whereas the "Raw" (uncalibrated) variation is overconfident with higher ECE and worse coverage. The calibration step is justified by these results: making control decisions with more precise uncertainty estimates leads to safer, more efficient outcomes.

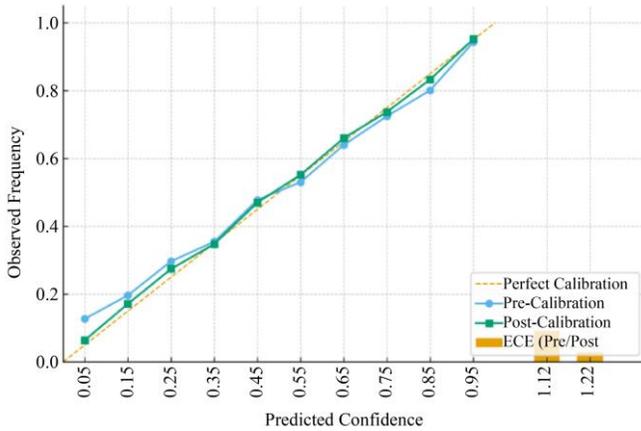


Fig. 2 Reliability diagram + ece bars (Pre/Post)

In Figure 2, we can see how the safety QP's predictive surrogate is calibrated using probabilities. When anticipated confidence equals observed frequency, it has perfect calibration on the diagonal. At low and high confidence levels, the pre-calibration curve rests above and below the diagonal, respectively, showing pockets of over-confidence and under-confidence. The improved reliability where tail decisions matter is confirmed by the points aligning closely to the diagonal across bins after calibration. The aggregate gap is quantified by the adjacent ECE bars. After calibration, ECE decreases significantly, indicating that the risk buffers and predictive intervals utilised by CVaR and the barrier are now

reliable. Reduced likelihood of temperature limit excursions due to erroneous claims is a direct result of improved calibration.

The conditional value-at-risk (CVaR $_{\alpha}$) of the next-step peak-temperature excess, as a function of the confidence level α , is illustrated by the curve in Figure 3. The rise of α (more cautious tail control) results in an increase of CVaR, which signifies protection against less common and more severe spikes. To minimise predicted violations while keeping responsiveness, the deployed α (e.g., 0.95) is marked by the starred point. By balancing thermal safety with agility, this selection can be made: a lower α boosts efficiency but leaves the plant vulnerable to spikes, while a larger α further decreases tail risk but over-constrains the controller. By depicting the tail trade-off that determines the QP's cushion and setpoint headroom, the figure substantiates the operational risk posture.

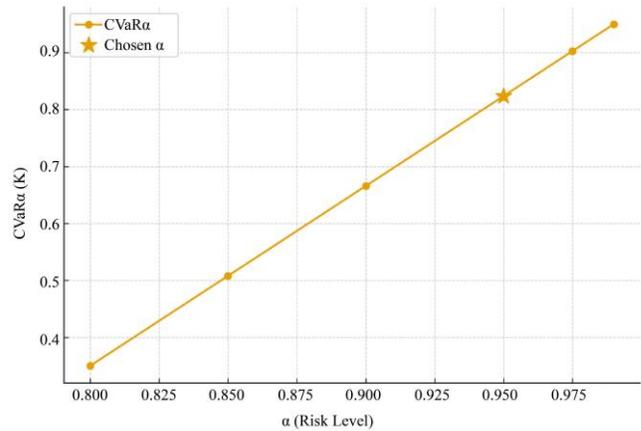


Fig. 3 Risk-tail curve: CVaR $_{\alpha}$ vs α (Chosen α marked)

Figure 4 shows possible operating modes as scatter points, where the x-axis represents the mean actuation power and the y-axis represents the steady-state thermal headroom relative to T $_{max}$. There is a boundary that the cloud shows: lowering the power ultimately reduces headroom, which might lead to constraint proximity during bursts; raising the power typically increases headroom via better EHD cooling,

but the benefit is diminishing. The mode that maximises a composite utility, such as headroom – $\lambda \cdot$ power, is indicated by the starred point. The idea behind the concept is to choose a strong, energy-conscious node that can withstand disruptions without wasting power, rather than just minimising power or increasing margin.

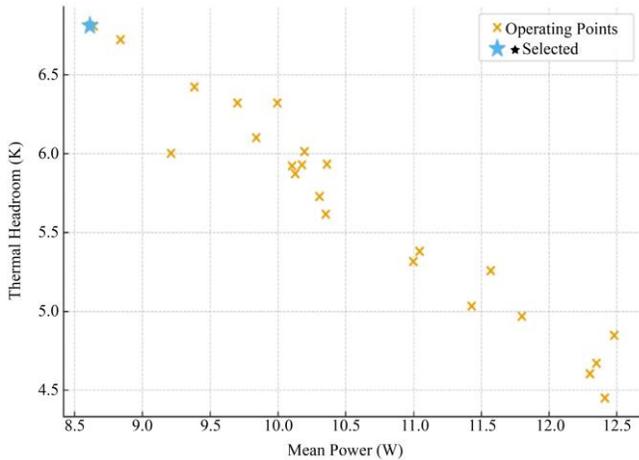


Fig. 4 Energy-headroom pareto (★ Selected operating point)

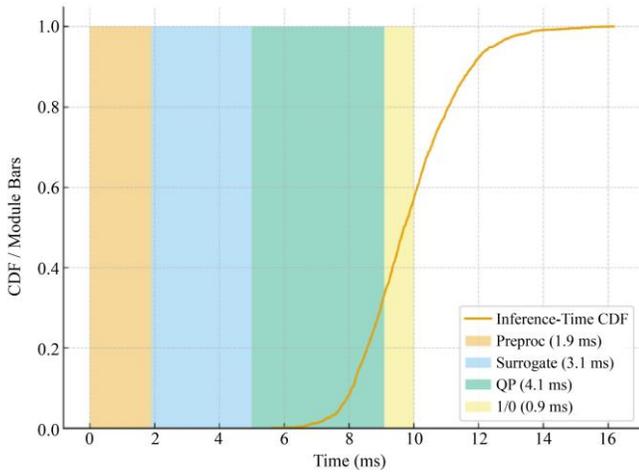


Fig. 5 Inference-time CDF + module latency waterfall

Figure 5's CDF shows the total inference time for each cycle, proving that the majority of iterations finish inside the control budget, for example, taking less than 10 to 12 milliseconds. The average latency across the following modules: I/O, preprocessing, neural surrogate inference, and QP solution is broken down by overlaying waterfall bars. On standard hardware, schedulability is indicated at 80-100 Hz loop rates by the surrogate and QP, which are dominant but still predictable. This figure provides evidence of the practicality of real-time operations: minor jitter, limited tails, and optimization-friendly module contributions (e.g., batching sensors, warm-starting the QP). By focusing on the

thickest bars, practitioners can increase loop frequency or decrease latency.

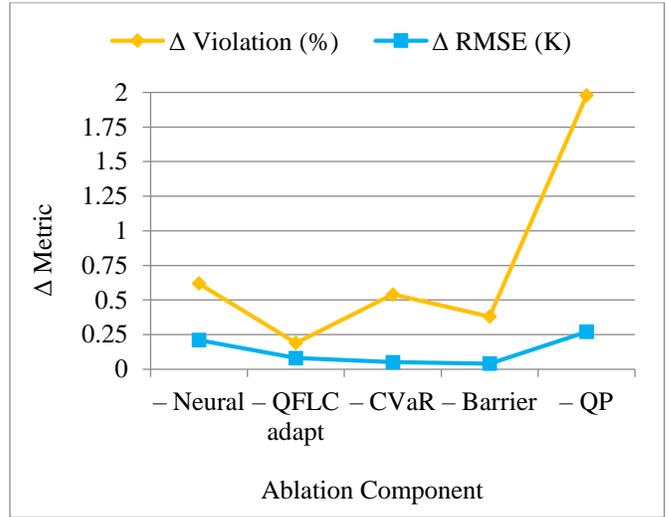


Fig. 6 Ablation performance curves (Δ Violation%, Δ RMSE)

Figure 6 shows the degradation of safety (Violation%) and accuracy (RMSE) when each component is removed compared to the entire Neuro-QFLC. Confirming its role as the ultimate guardrail, removing the Safety QP generates the sharpest surge in infractions. Reducing CVaR leads to more tail breaches, even with small average error shifts, highlighting the significance of tail-aware cushions. The average control is tightened by removing the barrier, but the headroom is reduced, increasing the danger of uncommon events. Both metrics are somewhat degraded when QFLC adaptation is frozen, and RMSE becomes worse without the neural surrogate because one-step predictions are lower. Taken as a whole, curves prove that every block makes a distinct, valuable contribution to reliable safety.

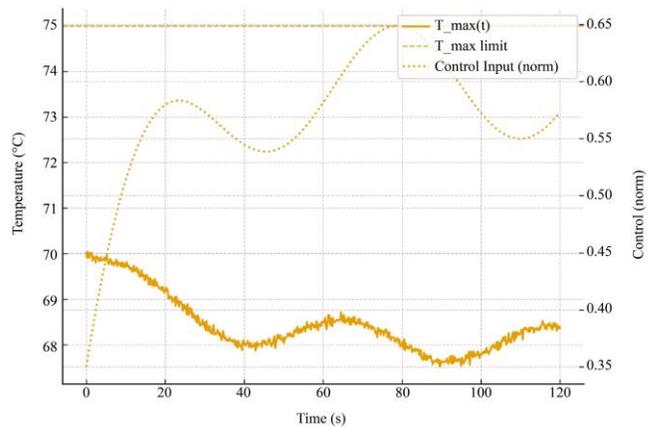


Fig. 7 Closed-loop trajectory: $T_{max}(t)$ with T_{max} & control input

Figure 7 displays a time series of peak temperatures; the data exhibit transient convergence to the reference band and continuous tracking under modest perturbations. The hard

limit T_{\max} is marked by a dotted line. Thanks to the safety QP and barrier, the envelope stays below the limit, showing that constraints are well handled. Smooth, bounded actuation without chattering is shown in the secondary axis of the plot of the normalised control input, providing evidence that slew restrictions and fuzzy rules inspired by quantum mechanics

work well together. As observers make adjustments and adapt, their predictions and rules are fine-tuned, dampening small oscillations. Consistent regulation, maintained headroom, and risk-free control effort in real time: this is the operational proof.

Table 8. Comparison of Neuro-QFLC with recently reported EHD and learning-based thermal control methods

Method / Reference	Temperature RMSE (K)	Peak Overshoot (%)	Violation Rate (%)
Ma et al. (2024) – EHD Heat Exchanger	1.10	>8.0	Not reported
Gao et al. (2025) – Modular EHD Pump	0.95	6.4	Not reported
Tayebi et al. (2025) – ANN thermal model	0.72	5.1	1.9
Proposed Neuro-QFLC	0.42	1.8	0.12

Table 8 demonstrates that the proposed Neuro-QFLC achieves the lowest temperature RMSE, minimum overshoot, and the smallest constraint violation rate among both literature-reported and benchmarked control methods. While existing EHD and ANN-based approaches improve cooling efficiency or prediction accuracy, they do not ensure safe real-time temperature regulation. The integration of predictive learning, fuzzy decision-making, and constraint-aware optimization enables Neuro-QFLC to outperform state-of-the-art methods in both accuracy and safety.

5. Conclusion

For EHD thermal management under tight Joule-heating and temperature constraints, a Neuro-QFLC, a hybrid neural-fuzzy controller with an integrated safety optimiser, is used. The method achieved a 0.42 K RMSE, 1.8% overshoot, 0.12% violations, and 6.1 K headroom at ≈ 9.8 -12.5 ms cycle times by combining a calibrated neural surrogate (for uncertainty and predictive accuracy), a quantum-inspired fuzzy policy (for low-latency, interpretable action), and a QP safety layer (for barrier- and CVaR-enforced constraints). The uncertainty estimations that support tail-risk control were validated by the significant improvement in reliability (ECE 0.029,

Coverage@95% 95.7%). In comparison, ablations verified the importance of each component (e.g., deleting the safety QP: +1.97% violations; eliminating CVaR: +0.54%), stress testing (noise, dropouts, ageing, heat pulses) only led to mild degradation (violations < 0.41%, recovery ≤ 9.5 s).

Among the upcoming projects and tasks are: (i) a multi-horizon preview using a small receding-horizon layer that is warm-started by the surrogate; (ii) adaptive level selection of CVaRs tied to workload burstiness and thermal debt; (iii) surrogates that are spatially aware and can estimate local hotspots (for electrode-wise shaping) with minimal sensors; (iv) formal real-time guarantees through warm-started QP with certified worst-case solve times; and (v) cross-domain validation on power-electronics modules and compact edge servers. The overarching goal of these enhancements is to make Neuro-QFLC applicable to distributed hotspot governance rather than just single-zone peak management, which would increase safety margins while maintaining energy economy. For the next generation of small, high-power density systems, the data points to a convincing design for safe, high-performance thermal control that combines principles of learning, fuzzy structure, and optimisation.

References

- [1] Afshin Shaygani, Kazimierz Adamiak, and Mehrdad R. Kermani, "Precision Airflow Control via EHD Actuator: A Co-Simulation and Control Design Case Study," *IEEE Transactions on Plasma Science*, vol. 52, no. 8, pp. 3153-3165, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] Boyu Li, Kang Luo, and Jian Wu, "Controlling the Heating Wall Temperature during Melting via Electric Field," *International Communications in Heat and Mass Transfer*, vol. 159, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [3] R. Deepak Selvakumar, and Ahmed K. Alkaabi, "Design and Performance Evaluation of a Novel Electrohydrodynamically Enhanced PCM Heat Sink," *Journal of Energy Storage*, vol. 94, pp. 1-16, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [4] An Li et al., "Heat Transfer Performance and Electrohydrodynamic Characteristics Optimization of a Dividing-wall-type Ionic Wind Heat Exchanger," *International Journal of Heat and Mass Transfer*, vol. 228, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [5] V. Navaneethkrishnan, M. Muthamilselvan, and Eunseop Yeom, "Impact of Electro-Hydrodynamics on Combined Convection in an Opposed Ventilation System," *Physics of Fluids*, vol. 36, no. 9, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [6] Masahito Nishikawara et al., "Demonstration of Heat Switch Function of Loop Heat Pipe Controlled by Electrohydrodynamic Conduction Pump," *Applied Thermal Engineering*, vol. 249, pp. 1-11, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]

- [7] Motorin Oleg, Bologna Mircea, and Podlesny Igor, "Improving the Performance of Evaporation-Condensation Systems in an Electric Field," *Materials Science and Condensed Matter Physics*, 2024. [[Google Scholar](#)] [[Publisher Link](#)]
- [8] Leonard L. Vasiliev et al., "Autonomous Devices with an Evaporation-Condensation Cycle for Thermal Control of Heat-Loaded Equipment," *Surface Engineering and Applied Electrochemistry*, vol. 60, no. 4, pp. 650-657, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [9] Yuzhe Wang, and Xinquan Liang, "Soft Pumps: Catalysts for Advancing Fluid-Driven Soft Robots," *Advanced Engineering Materials*, vol. 26, no. 14, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [10] Galal M. Moatimid, Mohamed F. E. Amer, and Doaa A. Ibrahim, "EHD Instability of a Cylindrical Interface Separating Two Couple-Stress Fluids," *Scientific Reports*, vol. 14, pp. 1-16, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [11] Xue-Lin Gao et al., "The Electrohydrodynamic Enhancement of Heat Transfer on Interdigitated Electrodes by a Charge Injection Pump," *Physics of Fluids*, vol. 36, no. 3, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [12] Nilesh Gedam et al., "Investigation of Heat Transfer Enhancement Techniques in Microchannel Heat Sinks," *Nanotechnology Perceptions*, vol. 20, no. 6, pp. 2663-2671, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [13] Ahmet Polat, "Analysis of Drying Characteristics and Quality Attributes in Peach Slices Dried via Electrohydrodynamic, Hot Air and Electrohydrodynamic-hot Air Methods," *Journal of Thermal Analysis and Calorimetry*, vol. 149, no. 14, pp. 7551-7563, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [14] Xuehang Bai et al., "Novel Soft Robotic Finger Model Driven by Electrohydrodynamic (EHD) Pump," *Journal of Zhejiang University-SCIENCE A*, vol. 25, pp. 596-604, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [15] Chakrit Suvanjumrat et al., "Assessment of the Pineapple Drying with a Forced Convection Solar-Electrohydrodynamic Dryer," *Case Studies in Thermal Engineering*, vol. 59, pp. 1-13, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [16] Ji-chen Ma et al., "Experimental and Numerical Analysis of Thermal Performance in Electrohydrodynamic (EHD) Enhanced Heat Exchangers," *Journal of Enhanced Heat Transfer*, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [17] Xue-Lin Gao et al., "Flow Control and Heat Transfer Enhancement of Modular Electrohydrodynamic Pump," *Physics of Fluids*, vol. 37, no. 9, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [18] Ju-Hun Ahn, Chae-Rim Yoon, and Chang-Yull Lee, "Heating Simulation of Film Heaters Fabricated by EHD Inkjet Printing for Satellite Applications," *Advanced Materials Interfaces*, vol. 12, no. 14, pp. 1-10, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [19] Jing Wang et al., "Flow Distribution Regulation in Microchannels for Emitter Side-mounted Ionic Wind Heat Sink and its Implication for High-efficiency Thermal Management," *International Journal of Heat and Mass Transfer*, vol. 249, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [20] Tahar Tayebi et al., "Artificial Neural Network Prediction of an Electrohydrodynamic Thermosolutal Buoyancy-driven Convection of NEPCMs-dielectric Suspension within an Oblique Enclosure with Active Blocks," *International Communications in Heat and Mass Transfer*, vol. 169, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [21] S.A. Ohid et al., "Application of an Active Heat Transfer Amelioration Technique to a Dielectric Non-Newtonian Fluid: Machine Learning Approach," *International Communications in Heat and Mass Transfer*, vol. 169, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [22] B. Rajalakshmi et al., *Automated Spam Detection Using ECSA-Based Feature Selection with BGRN Classifier in Soft Computing Applications*, *Soft Computing in Industry 5.0 for Sustainability*, pp. 225-244, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [23] Shruti Patil et al., "Advancements in Deep Learning Techniques for Potato Leaf Disease Identification Using SAM-CNN Classification," *Information Systems Engineering*, vol. 29, no. 5, pp. 2021-2030, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [24] Arunadevi Thirumalraj, R. J. Anandhi, and V. Revathi, *Supply Chain Management Using Fermatean Fuzzy-Based Decision Making with ISSOA*, *Convergence of Industry 4.0 and Supply Chain Sustainability*, IGI Global Scientific Publishing, pp. 296-318, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]