4th International Conference on Electronics, Engineering, Computer Science and Applied Development (EESD 2026)

Article

Power Semiconductor Selection and Drive Circuit Design for Brushless DC Motors in Unmanned Aerial Vehicles

Yicheng Wang[1,*]

[1] Qingdao Middle School, Qingdao, China
* Correspondence: Yicheng Wang, Qingdao Middle School, Qingdao, China

Abstract: Brushless direct-current (BLDC) motors are widely used in unmanned aerial vehicle propulsion, where power-semiconductor selection directly affects inverter loss, junction temperature, and electrical operating margin. However, conventional rating- or figure-of-merit-based selection provides limited representation of variable motor loads, while purely data-driven approaches may overlook explicit electrical and thermal constraints. This study proposes PCTO-BLDC, a physics-constrained computational framework that combines synchronized current–speed feature encoding, electro-thermal loss modeling, a lightweight neural surrogate, feasibility aggregation, and joint optimization of power devices and driver parameters. The surrogate is trained using physics-derived semiconductor-loss and junction-temperature targets, while candidate solutions remain subject to deterministic voltage, current, and thermal limits. Across ten repeated experiments, PCTO-BLDC achieved an average semiconductor loss of 4.55 ± 0.19 W, a peak junction temperature of 80.7 ± 2.5 C, and a robust feasibility rate of 95.1 ± 2.1%. External evaluation on DUDU-BLDC 1.5 retained a feasibility rate of 89.7 ± 3.6%. These results indicate that integrating measured operating profiles with physics-constrained computational optimization can support reproducible and interpretable power-stage design under defined BLDC operating conditions.

Keywords: unmanned aerial vehicle; brushless DC motor; power semiconductor selection; electro-thermal modeling; physics-constrained optimization

Received: 29 June 2026
Revised: 21 August 2026
Accepted: 31 August 2026
Published: 05 September 2026

1. Introduction

Brushless direct-current (BLDC) motors are widely used in unmanned aerial vehicle (UAV) propulsion because of their high power density, fast torque response, compact structure, and efficient electronic commutation [1]. Their operating performance depends strongly on the electronic speed controller and particularly on the semiconductor devices used in the three-phase inverter [2]. Power MOSFETs must withstand rapidly varying current, DC-link voltage, switching transients, and thermal stress while maintaining acceptable conduction and switching losses [3]. For UAV propulsion, inappropriate device selection may increase inverter loss, device temperature, and electrical stress even when nominal voltage and current ratings are satisfied.

Conventional power-device selection generally relies on maximum voltage, current rating, on-state resistance, or simplified figures of merit. These methods are transparent and computationally inexpensive, but they provide limited representation of time-varying motor conditions. In practical BLDC drives, current magnitude, rotational speed, switching frequency, dead time, gate-drive characteristics, and semiconductor temperature interact nonlinearly. A device with low room-temperature resistance may therefore not minimize total loss when switching behavior, reverse recovery, output capacitance, and temperature dependence are considered together [4]. Analytical electro-

thermal models provide strong physical interpretability, but repeated evaluation becomes costly when many devices and driver settings are searched jointly. Data-driven computational models can accelerate this process, yet unconstrained prediction may generate solutions that violate electrical or thermal limits [5]. Existing approaches also tend to separate motor operating-data analysis from semiconductor-level constraints [6]. A gap therefore remains in integrating measured operating profiles, physics-based electro-thermal estimation, computational prediction, and constrained optimization within a unified design process.

To address this problem, this study proposes a Physics-Constrained Thermal Optimization framework for BLDC drives, termed PCTO-BLDC. The main research objectives and contributions are as follows: (1) Operating-profile-aware feature representation for electrical stress. Synchronized BLDC current and rotational-speed measurements are transformed into compact operating descriptors that represent load magnitude, peak stress, signal variability, and rotational state. This enables device evaluation to reflect measured motor behavior rather than nominal ratings alone. (2) Physics-constrained neural surrogate for electro-thermal prediction. Semiconductor characteristics and motor operating features are jointly encoded by a lightweight neural network. Physically calculated conduction loss, switching loss, and junction temperature are used as supervised training targets, while voltage, current, and temperature limits remain explicit engineering constraints. (3) Constraint-aware co-optimization of power devices and driver parameters. MOSFET selection is coupled with switching frequency, dead time, and gate-drive settings so that conduction and switching losses can be balanced under defined electrical and thermal limits. The optimization therefore considers component choice and driver configuration within the same design process.

The proposed technical route integrates synchronized signal processing, operating-state feature encoding, semiconductor-parameter representation, physics-derived label generation, neural surrogate learning, and constrained numerical optimization. The surrogate estimates electro-thermal responses for candidate configurations, while the optimization mechanism searches the device and driver design space subject to explicit feasibility constraints. Candidate solutions are then checked using analytical electrical and thermal models to preserve physical consistency.

From an academic perspective, the approach connects data-driven computation with interpretable semiconductor electro-thermal modeling and provides a structured method for handling coupled component and control decisions. From an engineering perspective, it supports transparent loss decomposition, electrical derating, junction-temperature checking, and robustness assessment under variable operating conditions. The framework provides a reproducible basis for UAV BLDC power-stage component selection and driver design without replacing manufacturer specifications or formal hardware qualification procedures.

2. Related Works

2.1. Power Device Selection for BLDC Drive Systems

Power-device selection for BLDC inverters has traditionally relied on voltage margin, current capability, on-state resistance, switching characteristics, and thermal ratings [7]. This approach remains attractive because the design rules are transparent, computationally inexpensive, and compatible with manufacturer specifications. In UAV propulsion systems, rule-based screening can rapidly exclude devices unable to tolerate expected electrical stress and therefore serves as an efficient first-stage design method.

However, rating-based selection usually represents operating conditions using fixed or nominal points. This simplification is insufficient for UAV BLDC drives, where current and speed vary with throttle commands, propeller loading, acceleration, and disturbances. Selecting a device mainly according to low on-state resistance may reduce conduction loss while increasing gate-drive requirements or switching loss. Simplified figures of merit partially capture this trade-off, but they compress multiple electrical characteristics into a

single scalar and provide limited information about individual loss mechanisms under different conditions [8]. Conventional screening is therefore useful for feasibility filtering but less suitable for final component ranking under variable loads.

2.2. Electro-Thermal Modeling and Driver Parameter Design

Physics-based electro-thermal modeling provides a more detailed alternative. Conduction loss, switching loss, reverse-recovery effects, capacitance-related energy, and thermal resistance can be combined to estimate semiconductor loss and junction temperature [9,10]. Its main advantage is interpretability because changes in device temperature or power dissipation can be traced to specific electrical mechanisms. Such models also allow voltage, current, and temperature limits to be imposed explicitly [11].

Nevertheless, model accuracy depends on parameter quality and assumptions. Datasheet values are commonly specified under limited test conditions, whereas practical switching behavior varies with temperature, current, bus voltage, gate resistance, and circuit parasitics [8,11]. Simplified steady-state thermal models may also underestimate transient temperature variation during rapid load changes [10]. Repeated electro-thermal calculations become inefficient when many MOSFET candidates and driver settings are evaluated jointly. Consequently, semiconductor selection and driver configuration are often treated separately despite their coupled influence on power loss.

Multi-objective and heuristic optimization can search larger discrete and continuous parameter spaces and identify trade-offs among efficiency, temperature, and electrical margin. Their main limitation is computational cost because each candidate may require repeated analytical evaluation. Results can also depend on predefined objective weights and search ranges, reducing reproducibility across implementations.

2.3. Data-Driven and Computational Approaches

Recent computational approaches have introduced machine learning into motor-state analysis, thermal prediction, semiconductor modeling, and converter optimization. Data-driven models can approximate nonlinear mappings more efficiently than repeated numerical calculations after training. Neural networks are particularly useful when operating signals and semiconductor parameters interact through relationships that are difficult to represent with a single closed-form equation [12,13]. Surrogate-assisted design can further reduce repeated simulation during converter parameter exploration [14,15].

However, purely data-driven methods introduce another limitation: prediction accuracy does not ensure physical feasibility. A model may estimate low loss for a candidate that violates voltage, current, or thermal limits when these constraints are not explicitly enforced [13,14]. In addition, many motor-oriented learning studies focus on fault classification or operating-state recognition rather than power-device selection. Conversely, semiconductor surrogate models often use device-level test variables without incorporating measured motor load profiles [12,13]. These research directions therefore remain relatively disconnected.

2.4. Research Gap and Positioning of This Study

Existing methods present a clear trade-off. Rule-based selection is transparent but insufficiently sensitive to dynamic operating conditions; electro-thermal modeling is interpretable but expensive for large searches; optimization methods improve parameter exploration but require repeated model evaluation; and data-driven approaches provide efficiency but may lack explicit physical constraints.

The remaining gap is therefore not the absence of loss models, optimization algorithms, or neural predictors, but their limited integration for UAV BLDC power-stage design. In particular, measured operating profiles are rarely connected directly with semiconductor electro-thermal prediction and simultaneous driver-parameter optimization.

This study addresses this gap by combining operating-profile-aware feature representation, physics-derived electro-thermal supervision, a lightweight neural surrogate, and explicit feasibility constraints within one computational framework. The

neural model accelerates candidate evaluation rather than replacing physical equations, while final device and driver configurations remain subject to analytical electrical and thermal verification. This design preserves interpretability and engineering constraints while reducing the computational burden of joint component and driver-parameter selection.

3. Methodology

3.1. Overall Architecture

This study proposes PCTO-BLDC, a physics-constrained computational framework that integrates measured BLDC operating data, semiconductor electro-thermal modeling, neural surrogate learning, and constrained power-stage optimization. The framework is designed as a hybrid physical–data-driven system rather than a purely neural predictor. Analytical models generate physically interpretable training targets and enforce electrical and thermal limits, while the neural surrogate accelerates repeated evaluation of candidate device–driver combinations (As shown in Figure 1).

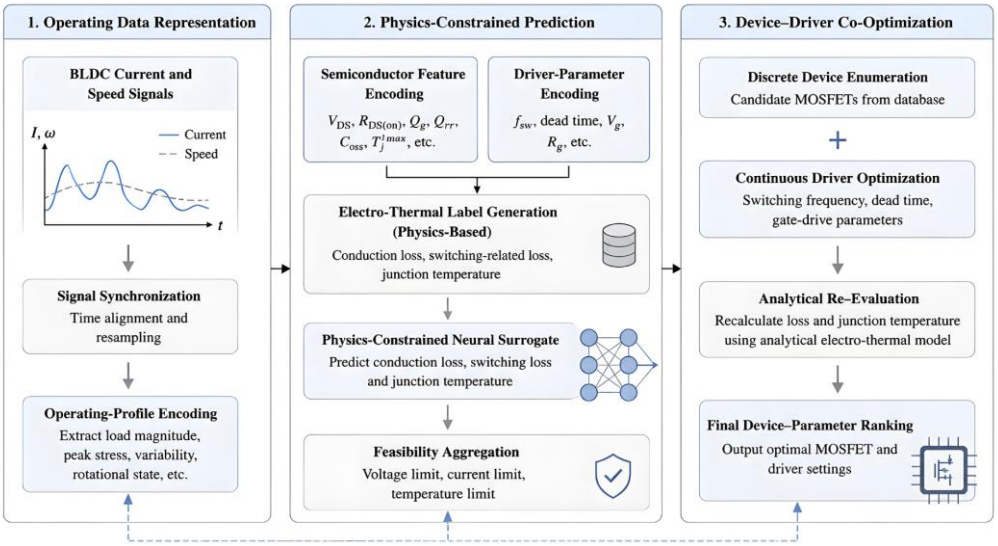

Figure 1. Overall Architecture of the Proposed Pcto-bldc Framework

The surrogate adopts a lightweight multilayer perceptron with hidden dimensions of 64, 32, and 16. Each hidden layer uses ReLU activation and layer normalization, while the output layer uses linear activation to predict total semiconductor loss and junction temperature. Continuous inputs are standardized using statistics calculated only from the training partition. The complete optimization follows an outer discrete search over candidate MOSFETs and an inner continuous update of driver parameters, which avoids applying gradient operations directly to discrete device identities.

3.2. Data Synchronization and Operating-Profile Encoding

Current and rotational-speed measurements must describe the same physical interval before feature extraction. When the two signals are not sampled at identical timestamps, rotational speed is interpolated onto the current-signal timeline. For current timestamp t_i , the synchronized speed is

$$\omega(t_a) + \frac{t_i - t_a}{t_b - t_a}[\omega(t_b) - \omega(t_a)], \quad t_a \le t_i \le t_b \qquad (1)$$

where $\tilde{\omega}(t_i)$ is the synchronized rotational speed, t_a and t_b are adjacent original speed timestamps, and $\omega(t_a)$ and $\omega(t_b)$ are the corresponding measured speeds.

Each synchronized window is then represented by

$$\mathbf{x}_o = \left[I_{\mathrm{rms}}, I_{\mathrm{pk}}, \bar{I}, \sigma_I, \bar{\omega}, \sigma_\omega, V_{\mathrm{dc}}\right] \qquad (2)$$

where I_{rms} and I_{pk} are RMS and peak current, $\bar{I}$ and σ_I are current mean and standard deviation, $\bar{\omega}$ and σ_ω are rotational-speed statistics, and V_{dc} is DC-link voltage.

Windows from the same physical recording are kept within the same data partition to avoid information leakage.

3.3. Temperature-Dependent Conduction-Loss Modeling

The conduction component is calculated from measured current and temperature-dependent on-state resistance:

$$P_{\mathrm{cond}} = I_{\mathrm{rms}}^2 R_{\mathrm{on},0}\left[1 + \alpha_R\left(T_j - T_0\right)\right]D \quad (3)$$

where P_{cond} is conduction loss, $R_{\mathrm{on},0}$ is on-state resistance at reference temperature T_0 , α_R is the resistance temperature coefficient, T_j is junction temperature, and D is effective conduction duty ratio. This formulation prevents device ranking from relying only on room-temperature resistance.

3.4. Switching, Recovery, and Dead-Time Loss

Switching-related mechanisms are evaluated jointly as

$$P_{\mathrm{sw}} = f_{\mathrm{sw}}\left[\frac{1}{2}V_{\mathrm{dc}}I_{\mathrm{sw}}\left(t_r + t_f\right) + E_{\mathrm{oss}} + V_{\mathrm{dc}}Q_{\mathrm{rr}} + V_F I_{\mathrm{dt}} t_d\right] \quad (4)$$

where P_{sw} is switching-related loss, f_{sw} is switching frequency, I_{sw} is switching current, t_r and t_f are effective transition times, E_{oss} is output-capacitance energy, Q_{rr} is reverse-recovery charge, V_F is body-diode forward voltage, I_{dt} is dead-time current, and t_d is dead time. The formulation captures the trade-off between reduced current ripple at higher switching frequency and increased switching energy.

3.5. Dynamic Electro-Thermal State Estimation

A first-order thermal network is used to propagate total device loss into junction temperature:

$$T_j^{(k+1)} = T_c + \left(T_j^{(k)} - T_c\right)e^{-\Delta t/\tau_{\mathrm{th}}} + P_{\mathrm{loss}}^{(k)} R_{\mathrm{th}}\left(1 - e^{-\Delta t/\tau_{\mathrm{th}}}\right) \quad (5)$$

where $T_j^{(k)}$ is junction temperature at update step k , T_c is case temperature, Δt is the update interval, τ_{th} is the effective thermal time constant, R_{th} is thermal resistance, and $P_{\mathrm{loss}}^{(k)}$ is total semiconductor loss. The resulting $T_j^{(k+1)}$ is fed back into Eq. (3), creating an iterative electro-thermal coupling rather than a fixed-temperature correction.

3.6. Physics-Constrained Neural Surrogate and Training Labels

The surrogate input concatenates the operating descriptor $\mathbf{x}_o$, semiconductor feature vector $\mathbf{x}_d$, and driver-control vector $\mathbf{u}$. The hidden representation and regression outputs are

$$\mathbf{h}^{(l)} = \mathrm{LN}\left(\mathrm{ReLU}\left(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right)\right), \quad l = 1,2,3, \qquad \hat{\mathbf{y}} = \mathbf{W}^{(4)}\mathbf{h}^{(3)} + \mathbf{b}^{(4)} \quad (6)$$

where $\mathbf{h}^{(0)}$ is the standardized concatenated input, $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ are trainable weights and biases, ReLU(·) is the rectified linear activation, LN(·) denotes layer normalization, and $\hat{P}_{\mathrm{loss}}$ and $\hat{T}_j$ are the predicted total loss and junction temperature.

Training labels are generated from Eqs. (3)–(5), rather than manually defined. The regression objective is

$$\mathcal{L}_{\mathrm{sur}} = \frac{1}{M}\sum_{m=1}^{M}\left[\left(\hat{P}_m - P_m\right)^2 + \lambda_T\left(\hat{T}_m - T_m\right)^2\right] \quad (7)$$

where M is the number of training samples, P_m and T_m are physics-derived labels, and λ_T balances the two regression targets. Hard electrical constraints are not learned by the network and remain deterministic checks in the subsequent optimization stage.

3.7. Explicit Feasibility Aggregation

Prediction accuracy alone is insufficient for power-device selection. Each device–driver combination must satisfy electrical and thermal limits over the evaluated operating windows. For window m , feasibility and the aggregated feasibility ratio are defined as

$$c_m(d,\mathbf{u}) = \mathbf{1}\left(V_m \le \gamma_V V_{\mathrm{DS,max}},\ I_m \le \gamma_I I_{\mathrm{D,max}},\ T_{j,m} \le T_{j,\max}\right), R_{\mathrm{feas}}(d,\mathbf{u}) = \frac{1}{M}\sum_{m=1}^{M} c_m\,(d,\mathbf{u}) \quad (8)$$

where d denotes a candidate device, $\mathbf{u}$ is the driver-parameter vector, γ_V and γ_I are voltage and current derating coefficients, and $V_{\mathrm{DS,max}}$, $I_{\mathrm{D,max}}$, and $T_{j,\max}$ are device limits. This mechanism separates prediction quality from engineering feasibility: a low-loss candidate cannot obtain the final rank if it repeatedly violates hard constraints.

3.8. Discrete–Continuous Co-Optimization and Parameter Update

The optimization uses an outer loop to enumerate discrete MOSFET candidates. For each fixed device d , the continuous driver parameters $\mathbf{u}$ are optimized independently. The objective and projected update are

$$J(d,\mathbf{u}) = w_P\tilde{P}_{\text{loss}} + w_T\tilde{T}_j + w_F(1-R_{\text{feas}}) + \Phi(d,\mathbf{u}) \quad (9)$$

$$\mathbf{u}^{(k+1)} = \Pi_\Omega\left(\mathbf{u}^{(k)} - \eta_k\nabla_{\mathbf{u}}J\left(d,\mathbf{u}^{(k)}\right)\right) \quad (10)$$

where $\tilde{P}_{\text{loss}}$ and $\tilde{T}_j$ are normalized loss and temperature terms, w_P , w_T , and w_F are objective weights, $\Phi(d,\mathbf{u})$ assigns a large penalty to hard-limit violations, η_k is the update step, and $\Pi\Omega(\cdot)$ projects the updated parameters into the admissible control domain Ω .

The discrete device identity is never differentiated. Instead, each candidate receives its own optimized continuous parameter set, after which all feasible device–parameter pairs are re-evaluated with the analytical model and ranked using the same objective. This outer-enumeration/inner-optimization structure makes the co-optimization mechanism computationally explicit and physically traceable.

Algorithm 1. PCTO-BLDC Device and Driver Co-Optimization

Input: BLDC dataset D, device set S, control domain Ω, constraint set C, maximum iteration number K

Output: selected device d*, optimized driver parameters u*, robust feasibility ratio R*

1: Synchronize current and speed signals in D using Eq. (1)
2: Group samples by physical recording before window extraction
3: Construct operating descriptors using Eq. (2)
4: Encode semiconductor features for every device d in S
5: Generate conduction-loss labels using Eq. (3)
6: Generate switching-loss labels using Eq. (4)
7: Update thermal labels recursively using Eq. (5)
8: Standardize training inputs and train the surrogate using Eqs. (6)–(7)
9: For each discrete device d in S, initialize driver parameters u in Ω
10: Predict loss and temperature for the current device–parameter pair
11: Compute hard-constraint feasibility using Eq. (8)
12: Evaluate the constrained objective J using Eq. (9)-(10)
13: Update continuous driver parameters and project them back to Ω
14: Re-evaluate all shortlisted pairs with the analytical model
15: Return d*, u*, and the corresponding robust feasibility ratio R*

The proposed methodology therefore forms a complete computational loop from synchronized BLDC measurements to physically supervised neural prediction and constrained device–driver optimization. Its main distinction lies in combining data-driven acceleration with explicit semiconductor physics: the surrogate reduces repeated electro-thermal evaluation, whereas device ratings, thermal limits, and final candidate verification remain analytical and directly interpretable.

4. Results and Analysis

4.1. Experimental Setup

Experiments were conducted on a workstation with an Intel Core i7-12700K CPU, 32 GB RAM, and an NVIDIA RTX 3060 GPU using Python 3.11 and PyTorch. The PCTO-BLDC surrogate used the 64–32–16 architecture, Adam optimization, an initial learning rate of 0.001, batch size of 128, and a maximum of 300 epochs. Early stopping was activated after 25 epochs without validation improvement. All experiments were independently repeated ten times ($n = 10$) using different grouped splits and random seeds.

Two public BLDC datasets were used. The primary DUDU-BLDC dataset contains synchronized current and rotational-speed measurements from Fein EC48 motors under healthy, mechanical-imbalance, electrical-degradation, and combined conditions. Signals were sampled at 50 kHz and segmented into 0.8-s windows, with 28 statistical and spectral features. The dataset is publicly available under CC BY 4.0. DUDU-BLDC 1.5 was used

only for external evaluation. This independent 2026 acquisition includes five motors and 50 recordings sampled at 20 kHz. After preprocessing, eight 0.8-s windows were retained per recording, giving 400 windows with 26 engineered features. Raw data, processed data, metadata, and preprocessing code are released under CC BY 4.0.

Current and speed channels were synchronized before feature extraction. Continuous inputs were standardized using statistics calculated only from the training partition. Source recordings, rather than individual windows, were assigned to training, validation, and test subsets at 60:20:20 to prevent correlated windows from crossing partitions. DUDU-BLDC 1.5 was excluded from training and hyperparameter selection.

Four baselines were used: Rating-Margin Selection (RMS), figure-of-merit ranking (FOM), Analytical Weighted-Sum search (AWS), and NSGA-II electro-thermal optimization. The three primary metrics were average semiconductor loss, peak junction temperature, and robust feasibility rate (RFR), defined as the percentage of windows satisfying all electrical and thermal constraints. Results are reported as mean ± standard deviation. Paired t -tests or Wilcoxon signed-rank tests were applied according to normality, with Holm correction and $p < 0.05$ as the significance threshold.

4.2. Performance Comparison

Table 1 compares PCTO-BLDC with the four baselines. The proposed framework achieved the lowest average semiconductor loss of 4.55 ± 0.19 W and peak junction temperature of 80.7 ± 2.5 C, with an RFR of 95.1 ± 2.1%.

Table 1. Overall Performance Comparison ($n = 10$).

Method	Semiconductor Loss (W) ↓	Peak Junction Temperature (°C) ↓	RFR (%) ↑
RMS	5.84 ± 0.31	91.6 ± 3.8	82.7 ± 4.9
FOM	5.37 ± 0.28	89.4 ± 4.1	84.6 ± 4.2
AWS	4.98 ± 0.24	84.8 ± 3.2	90.3 ± 3.1
NSGA-II	4.62 ± 0.21	82.9 ± 2.8	93.8 ± 2.7
PCTO-BLDC	4.55 ± 0.19	80.7 ± 2.5	95.1 ± 2.1

Compared with RMS and FOM, PCTO-BLDC improved all three metrics significantly after Holm correction ($p < 0.01$). Relative to AWS, reductions in loss and peak temperature remained significant ($p = 0.008$ and $p = 0.012$). Against NSGA-II, the loss difference was not significant ($p = 0.184$), while peak temperature remained lower ($p = 0.041$); the RFR difference was also not significant ($p = 0.073$). The results therefore indicate a more favorable electro-thermal balance rather than uniform superiority over the strongest baseline.

This behavior is physically interpretable. Low-resistance devices favored by RMS or FOM reduced conduction loss in some high-current windows but were less advantageous after switching, recovery, and thermal feedback were considered. Final analytical verification prevented the neural surrogate from determining device ranking independently of engineering constraints.

4.3. Ablation and Mechanism Verification

Table 2 evaluates the principal mechanisms of PCTO-BLDC.

Table 2. Ablation Results ($n = 10$).

Configuration	Semiconductor Loss (W) ↓	Peak Junction Temperature (°C) ↓	RFR (%) ↑
Full PCTO-BLDC	4.55 ± 0.19	80.7 ± 2.5	95.1 ± 2.1
w/o operating-profile encoding	4.72 ± 0.25	82.4 ± 3.0	90.8 ± 3.5
w/o thermal constraint	4.43 ± 0.20	88.9 ± 3.7	86.4 ± 4.2
w/o joint driver optimization	4.91 ± 0.23	83.6 ± 3.1	92.2 ± 2.8
w/o feasibility aggregation	4.49 ± 0.22	86.2 ± 3.9	88.1 ± 4.6

Removing operating-profile encoding reduced RFR by 4.3 percentage points ($p = 0.012$), showing that nominal loading does not preserve peak-stress information. Removing the thermal constraint slightly reduced average loss but increased peak temperature by 8.2 C and lowered RFR ($p < 0.001$), demonstrating that loss minimization alone does not ensure thermal robustness. Removing joint driver optimization increased loss by 0.36 W ($p = 0.006$). Without feasibility aggregation, loss again decreased slightly, but temperature and RFR deteriorated, confirming the need to exclude physically unfavorable candidates.

4.4. Convergence Analysis

Figure 2 shows optimization trajectories from ten independent runs with 95% confidence intervals. The normalized PCTO-BLDC objective decreased from 1.000 ± 0.041 to 0.541 ± 0.026 and stabilized after approximately 42 iterations. NSGA-II reached 0.567 ± 0.032 after 60 iterations. The final difference was significant ($p=0.038$), although the confidence intervals overlapped during several intermediate iterations.

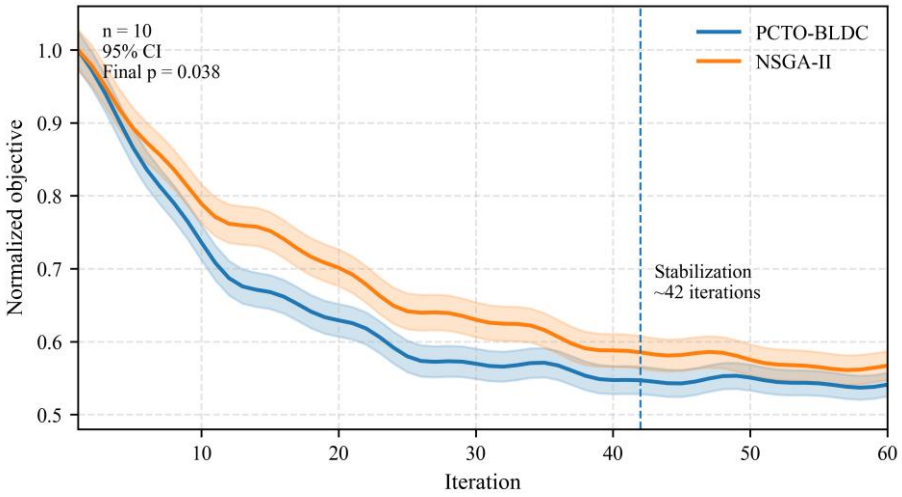

Figure 2. Convergence of the Normalized Optimization Objective over Ten Independent Runs; Shaded Regions Denote 95% Confidence Intervals.

The earlier stabilization is attributable to surrogate-assisted evaluation of continuous driver settings for each discrete device candidate. Because convergence was measured using the constrained objective rather than prediction loss alone, the curve reflects both electro-thermal performance and feasibility.

4.5. Stability, Generalization, and Trustworthiness Analysis

Figure 3 evaluates robustness under altered measurement and electrical conditions. Clean-data RFR was 95.1 ± 2.1%. With 2% current noise, RFR remained 94.4 ± 2.4% ($p = 0.318$). At 5% noise, it decreased to 92.3 ± 2.8% ($p = 0.027$). A ±10% DC-link voltage perturbation produced 91.6 ± 3.1% ($p=0.015$), while switching-frequency perturbation yielded 93.0 ± 2.5% ($p = 0.044$). External testing on DUDU-BLDC 1.5 resulted in 89.7 ± 3.6% ($p = 0.006$).

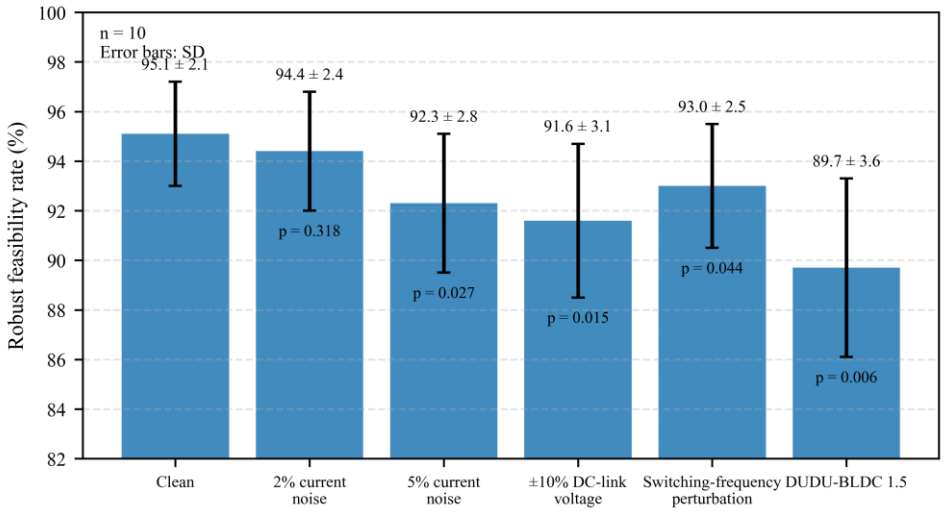

Figure 3. Robust Feasibility Rate under Measurement, Electrical, Control, and External-Dataset Perturbations ($n = 10$); Error Bars Indicate Standard Deviation.

The non-monotonic degradation reflects different physical mechanisms. Voltage variation directly changes switching stress, whereas frequency perturbation primarily shifts the balance between conduction and switching losses. The larger decrease on DUDU-BLDC 1.5 is consistent with its independent acquisition configuration and different motor population and is interpreted as cross-acquisition robustness rather than unrestricted fault generalization.

Trustworthiness was maintained by analytically re-evaluating every shortlisted configuration. Neural predictions accelerated candidate screening, while voltage, current, temperature, and datasheet limits remained deterministic constraints. This separation supports interpretability, robustness assessment, and traceable engineering compliance without implying formal UAV hardware certification.

5. Conclusion

The results show that incorporating measured operating profiles into UAV BLDC power-device evaluation improves the representation of electrical stress compared with nominal-condition selection. Removing the operating-profile encoding reduced the robust feasibility rate from 95.1 ± 2.1% to 90.8 ± 3.5% ($p = 0.012$), indicating that synchronized current and rotational-speed information contributes mainly to preserving peak-stress and variable-load characteristics rather than simply reducing average semiconductor loss.

The physics-constrained electro-thermal mechanism also produced a measurable effect on device selection. PCTO-BLDC achieved an average semiconductor loss of 4.55 ± 0.19 W and a peak junction temperature of 80.7 ± 2.5 C. When the thermal constraint was removed, average loss decreased slightly to 4.43 ± 0.20 W, but peak junction temperature increased to 88.9 ± 3.7 C and the feasibility rate fell to 86.4 ± 4.2% ($p < 0.001$). This result demonstrates that minimum predicted loss alone is not an adequate selection criterion and that explicit thermal limits are necessary for balanced device ranking.

The joint optimization of the power device and driver parameters further improved the electro-thermal operating point. Removing joint driver optimization increased semiconductor loss from 4.55 ± 0.19 W to 4.91 ± 0.23 W ($p = 0.006$). The complete framework also converged to a normalized objective of 0.541 ± 0.026 after approximately 42 iterations and retained an RFR of 89.7 ± 3.6% on the independent DUDU-BLDC 1.5 dataset. These findings support the use of coupled component and driver-parameter optimization rather than treating them as independent design stages.

Several limitations remain. The public datasets contain a limited number of physical BLDC motors and were not collected from airborne UAV propulsion systems. The electro-thermal model also simplifies package, PCB parasitics, transient thermal impedance, and

device-parameter uncertainty. Future work should therefore validate the method on UAV propulsion benches, incorporate transient thermal and parasitic models, and extend the candidate space to GaN and SiC devices under broader voltage and load conditions.

References

1. F. Saemi and M. Benedict, "Brushless DC Motor Sizing Algorithm for Small UAS Conceptual Designers," *Aerospace*, vol. 11, no. 8, p. 649, 2024.
2. F. Saemi, A. Whitson, and M. Benedict, "Heat transfer models and measurements of brushless dc motors for small uass," *Aerospace*, vol. 11, no. 5, p. 401, 2024.
3. H. E. Gündoğdu, F. Yabansu, F. O. Hocaoğlu, and T. Özer, "Drone'lar için Sensörsüz BLDC Motor Sürücü Tasarımı ve Uygulaması," *Afyon Kocatepe Üniversitesi Fen Ve Mühendislik Bilimleri Dergisi*, vol. 24, no. 5, pp. 1251–1259, 2024.
4. C. C. Hsu, Y. H. Hsieh, H. C. Tang, Y. M. Meng, M. S. Huang, P. H. Tseng, H. H. Lu, and H. Y. Hsu, "Bidirectional reduced-order electrothermal modeling of power MOSFETs for electric vehicle thermal management applications," *Case Studies in Thermal Engineering*, vol. 74, p. 106790, 2025.
5. D. A. Andrioaia and V. G. Gaitan, "Finding fault types of BLDC motors within UAVs using machine learning techniques," *Heliyon*, vol. 10, no. 9, 2024.
6. S. Yang, L. Zhao, X. Gu, and W. T. Ng, "An artificial neural network based electro-thermal behavioral model for SiC MOSFETs operating under wide temperature range and short-circuit condition," *Microelectronics Journal*, vol. 165, p. 106836, 2025.
7. A. Bahadir, O. Aydogdu, and E. Bahadir, "Three-Phase 75 kW Brushless Direct Current Motor for Electric Vehicles: Different Power Stage Design, Calculation of Losses, Cooling Techniques, and Comparison," *Applied Sciences*, vol. 14, no. 4, p. 1365, 2024.
8. V. Barba, S. Musumeci, F. Stella, F. Mandrile, and M. Palma, "Investigation of dead time losses in inverter switching leg operation: GaN FET vs. MOSFET comparison," *Energies*, vol. 17, no. 15, p. 3855, 2024.
9. X. Lan and N. Zhao, "Development of a steady state electrothermal cosimulation model of SiC power modules," *International Journal of Heat and Mass Transfer*, vol. 226, p. 125460, 2024.
10. C. Morel and J. Y. Morel, "Power semiconductor junction temperature and lifetime estimations: A review," *Energies*, vol. 17, no. 18, p. 4589, 2024.
11. Y. Zhao, Z. Wang, J. Wang, Y. Tang, B. Ji, C. Chen, and G. Li, "Electrical-thermal coupling modeling of SiC MOSFETs based on field-circuit coupling and its application in junction temperature calculation during surges," *IEEE Transactions on Power Electronics*, vol. 40, no. 2, pp. 2651–2667, 2024.
12. R. Wang, Y. Chen, S. Tong, C. Cheng, and Y. Kang, "Compensated neural network training algorithm with minimized training dataset for modeling the switching transients of SiC MOSFETs," *Energies*, vol. 17, no. 23, p. 6061, 2024.
13. G. Avila Saccol, V. H. Bui, and W. Su, "A deep-neural-network-based surrogate model for DC/AC converter topology selection using multi-domain simulations," *Energies*, vol. 17, no. 24, p. 6467, 2024.
14. A. Mahajan, S. Chaturvedi, S. Das, W. Su, and V. H. Bui, "Optimal parameter design for power electronic converters using a probabilistic learning-based stochastic surrogate model," *Next Energy*, vol. 9, p. 100464, 2025.
15. Y. Liu, P. Jayathurathnage, J. Kyyrä, and Y. Zhang, "Data-driven design of MHz resonant converters for full-freedom soft-switching optimization," *IEEE Transactions on Power Electronics*, vol. 41, no. 1, pp. 729–742, 2026.