electronics

MDPI

Article

Verifier-Gated Multi-Candidate Supervisory Selection for Constraint-Aware Flux-Weakening in Permanent Magnet Synchronous Motor Drives

Taehyun Kang and Han Ho Choi *

Division of Electronics and Electrical Engineering, Dongguk University-Seoul, 30, Pildong-ro 1-gil, Jung-gu, Seoul 04620, Republic of Korea
* Correspondence: hhchoi@dongguk.edu; Tel.: +82-2-2260-3777

Abstract

Flux-weakening current-reference generation must balance torque tracking, current use, voltage reserve, and continuity. This paper develops a verifier-gated multi-candidate supervisor for permanent magnet synchronous motor drives. Six deterministic reference laws are screened before a history-dependent score selects an admissible candidate. For a nonempty verified set, the encoded nominal-model constraints hold independently of ranking weights. Offline cross-entropy method calibration leaves the verifier unchanged. Across five quasi-steady profiles, the supervisor reduces mean squared current by 3.46% relative to a margin-matched scalar comparator, while torque root mean square error increases by 4.4%. Equal-budget retuning preserves this trade-off. An additional continuous nonlinear-programming comparator achieves lower current use and torque error than the supervisor, with higher measured host latency. Candidate ablations expose redundant options on the chosen profiles. Fixed-reference vertex checks satisfy the voltage limit for a synthetic $\pm 5\%$ parameter box, whereas larger boxes expose margin failures. Sixty sampled electrical closed-loop simulations include switching inverters, delays, and modeled nonidealities; no current-limit violation occurs in that panel. These results establish ranking-independent model-side admissibility and quantify its operating limits and computational trade-offs without implying physical-drive or real-time certification.

Keywords: permanent magnet synchronous motor; PMSM; flux weakening; current-reference generation; constraint verification; multi-candidate supervisory control; cross-entropy method; MTPA; MTPV

Check for updates

Academic Editor: Herminio Martínez-García

Received: 21 September 2026
Revised: 20 October 2026
Accepted: 22 October 2026
Published: 26 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Permanent magnet synchronous motors (PMSMs) are used in traction, servo, pump, compressor, fan, and appliance drives because they can provide high torque density and high efficiency [1–3]. In field-oriented control (FOC), electromagnetic torque is regulated through direct- and quadrature-axis current references. At low and medium speed, the reference is commonly selected along a maximum-torque-per-ampere (MTPA) trajectory. Above base speed, the inverter-voltage constraint becomes active, and the MTPA point may cease to be feasible; negative direct-axis current is then introduced to reduce the effective air-gap flux [4–9]. Flux-weakening (FW) current-reference generation couples several constraints and objectives. For a commanded torque and electrical speed, the issued reference should remain inside the stator-current and inverter-voltage limits, avoid excessive inter-update

changes, and provide acceptable torque tracking. A more negative direct-axis current can increase voltage headroom but also increase current magnitude and reduce the available quadrature-axis current. Conversely, operation near an estimated voltage boundary can improve torque tracking while increasing sensitivity to dc-link sag and parameter mismatch. Reviews of FW control distinguish calculation-based, voltage-loop, and predictive approaches and emphasize their different trade-offs in model dependence, dynamic response, calibration, and computation [10–12]. Classical current-vector and voltage-feedback methods handle wide-speed operation through predefined region-wise trajectories or feedback regulation [5,7–9,13–15]. Direct-torque, predictive, and analytical approaches instead use dynamic optimization or explicit reference calculation [16–22]. These method families supply complementary operating laws at different control layers; the present problem is how to coordinate several such laws while keeping admissibility independent of their ranking. Table 1 qualitatively positions these established method families, four closest recent studies, and the present work. The four closest recent studies address explicit MTPA-, MTPV-, maximum-current-, and loss-minimization-related approximations [23], scalar residual correction for one loss-minimization objective [24], speed- and voltage-margin scheduled two-anchor blending [25], and safeguarded continuous optimization of one registered convex performance index [26]. Because the cited studies use different machine models, control layers, and validation protocols, Table 1 compares decision architecture and constraint handling rather than numerical performance. The safeguarded scalar framework [26] is the closest technical antecedent.

Despite these advances, a distinct integration problem remains: heterogeneous current-reference laws must be coordinated with previous-reference and previous-option information, while current, voltage, branch, denominator, and inter-update admissibility remains invariant to the ranking weights. The proposed method addresses this problem by generating six deterministic candidates, constructing a verified set with a controller-model gate, and ranking only that set. CEM calibrates the score offline but is not required for the admissibility property.

Table 1. Qualitative positioning of representative PMSM FW and current-reference approaches and the proposed verifier-gated supervisor. The comparison concerns decision architecture and constraint handling rather than numerical performance because the cited studies use different machine models, control layers, and validation protocols.

Approach	Decision and Constraint Structure	Relation to the Present Study
Established method families		
Classical current-vector and region-based FW [5,7,8]	Uses predefined MTPA, constant-current/FW, and MTPV trajectories, with region transitions determined from current and voltage boundaries.	Provides the operating-region geometry and several candidate laws used here; the proposed architecture places heterogeneous laws in a common generation–verification–ranking pipeline.
Voltage-feedback and saturation-tracking FW [9,14,15]	Adjusts flux-producing current from voltage utilization, regulator saturation, or voltage-constraint error; behavior depends on feedback, filtering, and anti-windup design.	Provides online voltage-stress response through a feedback law; the proposed method instead verifies and ranks a finite set of independently generated references.
Direct-torque and predictive control [16–18]	Incorporates torque/flux dynamics and constraints through direct control, predictive control, or online dynamic optimization.	Operates at the dynamic lower-control layer. The present method is complementary: it supplies a verified current reference to the conventional inner FOC loop.
Analytical and MTPV-oriented reference generation [19–22]	Uses analytical, quartic, direct-flux, magnetic-model-dependent, or speed-limited formulations to generate selected current trajectories.	Supplies analytical candidate laws and benchmarks. The proposed method enables several such laws to coexist under one independent verifier and history-dependent selector.
Reviews and taxonomies [10–12]	Classifies calculation-based, voltage-loop, predictive, vector-current, and MTPV-oriented FW methods and summarizes implementation trade-offs.	Provides the taxonomy used to position the proposed supervisor across calculation-based, feedback, predictive, and MTPV-oriented approaches.

Table 1. *Cont.*

Approach	Decision and Constraint Structure	Relation to the Present Study
Closest recent studies and the present work		
Explicit approximations [23]	Generates one explicit MTPA-, MTPV-, maximum-current-, or loss-minimization-related trajectory at a time, with supplementary boundary logic.	Supplies low-cost candidate laws that can enter the finite set; the present study adds a common verifier and history-dependent ranking layer.
Residual-corrected LMC [24]	Applies one scalar loss-optimality correction to an explicit LMC warm start, with scalar feasibility or boundary management.	Supplies a residual-polished loss-minimizing reference that can serve as one candidate; the present study coordinates multiple laws and uses previous-option information.
SVS [25]	Continuously blends two anchors using speed and voltage-margin pressures, followed by projection and torque derating.	Supplies a continuous scheduled reference that can serve as one candidate; the present study adds finite-set verification and selection.
Safeguarded scalar framework [26]	Continuously minimizes one registered scalar objective over an objective-independent feasible interval with certified boundary search, derating, and deterministic fallback.	Closest antecedent for model-side feasibility and fallback. The present method adds discrete, history-dependent orchestration among heterogeneous candidate laws with ranking-independent admissibility.
Present study	Generates a finite candidate set; verifies current, voltage, branch, denominator, and inter-update conditions; and ranks only the verified candidates using previous-reference and previous-label information.	Separates heterogeneous candidate generation, weight-independent model-side admissibility, and history-dependent ranking.

This paper makes three contributions. First, it formulates a history-dependent finite-set supervisory architecture that combines heterogeneous reference laws, a conservative derated fallback, a reference-increment limit, and a mode-persistence term above the conventional inner current loop. Second, it separates candidate ranking from deterministic admissibility and proves that whenever the verified set is nonempty, any ranking rule restricted to that set returns a controller-model-admissible reference independently of its weights. Third, it separates performance trade-offs from verifier protection through a margin-matched scalar comparison, ranking-weight diagnostics, candidate ablations, grid refinement, parameter-envelope analysis, and sampled electrical closed-loop simulations.

The validation combines quasi-steady reference-generation tests with sampled electrical closed-loop simulation for one published IPMSM parameter set. The electrical simulations impose rotor speed and model current regulation, PWM, delay, and selected inverter and sensing effects. They do not represent a mechanical speed-control experiment or real-time hardware validation.

2. PMSM Model and Constrained Reference Problem

2.1. Steady-State Voltage and Torque Equations

In the synchronous d–q frame, the PMSM voltage equations are

$$v_d = R_s i_d + L_d \frac{di_d}{dt} - \omega_e L_q i_q, \tag{1}$$

$$v_q = R_s i_q + L_q \frac{di_q}{dt} + \omega_e (L_d i_d + \lambda_m), \tag{2}$$

where R_s is the stator resistance, L_d and L_q are the d- and q-axis inductances, λ_m is the permanent-magnet flux linkage, and ω_e is the electrical rotor speed. For quasi-steady current-reference calculation, the derivative terms are neglected:

$$v_d(i_d, i_q) = R_s i_d - \omega_e L_q i_q, \tag{3}$$

$$v_q(i_d, i_q) = R_s i_q + \omega_e (L_d i_d + \lambda_m). \tag{4}$$

The voltage magnitude is

$$v_s(i_d, i_q) = \sqrt{v_d(i_d, i_q)^2 + v_q(i_d, i_q)^2}. \tag{5}$$

The electromagnetic torque is

$$T_e(i_d, i_q) = \frac{3}{2} p\{\lambda_m + (L_d - L_q) i_d\} i_q, \tag{6}$$

where p is the number of pole pairs. The current and voltage constraints are

$$i_d^2 + i_q^2 \leq I_{\max}^2, \tag{7}$$

$$v_d(i_d, i_q)^2 + v_q(i_d, i_q)^2 \leq V_{\max}^2. \tag{8}$$

For space-vector PWM in the linear modulation range, $V_{\max}$ is approximated by $V_{\mathrm{dc}}/\sqrt{3}$ for the peak fundamental voltage-vector magnitude. Other modulation conventions can be accommodated by redefining $V_{\max}$.

All currents and voltages in the numerical study use peak-value synchronous-frame quantities. The source benchmark reports four poles and writes the torque coefficient as $3p/4$ with p denoting the pole count; Equation (6) instead uses pole pairs. The corresponding value used here is therefore $p = 2$ pole pairs. This convention makes (6) algebraically identical to the source benchmark torque equation.

2.2. Torque-Eliminated Scalar Representation

When the torque command is achievable, the torque equation can be used to write i_q as a function of i_d:

$$i_q(i_d; T_e^\star) = \frac{T_e^\star}{\frac{3}{2} p\{\lambda_m + (L_d - L_q) i_d\}}. \tag{9}$$

The current and voltage residuals along the torque curve are then

$$\Phi_I(i_d) = i_d^2 + i_q(i_d)^2 - I_{\max}^2, \tag{10}$$

$$\Phi_V(i_d) = v_d(i_d, i_q(i_d))^2 + v_q(i_d, i_q(i_d))^2 - V_{\max}^2. \tag{11}$$

The feasible set for exact torque tracking is

$$\mathcal{F}(T_e^\star, \omega_e) = \{(i_d, i_q(i_d)) : \Phi_I(i_d) \leq 0,\ \Phi_V(i_d) \leq 0\}. \tag{12}$$

For the motoring branch evaluated here, candidate generation is restricted to $T_e^\star \geq 0$, $-I_{\max} \leq i_d \leq 0$, and $i_q \geq 0$. Equation (9) is evaluated only where $|\lambda_m + (L_d - L_q) i_d| \geq 10^{-8}$ Wb. The primary reference-generation comparison uses 801 uniformly spaced points on $[-10, 0]$ A, giving $\Delta i_d = 0.0125$ A; the resolution study varies this number while keeping the laws and weights fixed. Ties are resolved by the first minimum on this ascending grid. These restrictions avoid the singular denominator in (9) and prevent an unintended generating-quadrant candidate. The set in (12) is a quasi-steady electrical feasibility set; it is not a closed-loop stability or transient-feasibility certificate.

In the proposed framework, exact torque tracking is preferred but not forced in all cases. If no exact-torque candidate satisfies the current and voltage constraints, a derated fallback candidate is allowed. This convention reflects physical operation: when the command exceeds the voltage-current capability, constraint compliance takes precedence over exact torque production.

3. Motivation for Verifier-Gated Multi-Candidate Supervision

Single-law and fixed-objective methods solve important but distinct FW reference-generation problems. Region-wise MTPA/FW rules are transparent and efficient, but their selected trajectory depends on a predefined boundary and transition rule. Voltage-feedback methods respond to voltage utilization online, with behavior governed by gain, filtering, saturation, and inner-loop dynamics. Projection and scalar-optimization methods explicitly enforce a voltage boundary or minimize a registered objective, but each produces one reference according to one decision law.

The safeguarded-scalar framework [26] provides the principal continuous-optimization benchmark used here: it constructs an objective-independent scalar feasible interval, derates infeasible torque commands, and minimizes a registered performance index with deterministic numerical safeguards. The present paper builds on that feasibility perspective but addresses a different integration layer. A drive designer may retain an MTPA point, a nominal projection, a torque-preserving margin point, a voltage-minimizing point, a residual-corrected point, and a conservative derated point because each expresses a distinct engineering intent.

The proposed supervisor treats these laws as a finite candidate set. Candidate-specific checks remain fixed, whereas the ranking can depend on torque error, voltage utilization, current magnitude, the previous reference, and the previous option label. This separation yields the principal architectural property: retuning, replacing, or misspecifying the ranking rule cannot authorize a candidate outside the verified set. The comparison with the margin-matched scalar J_ρ method therefore quantifies the current–torque–margin trade-off between continuous fixed-objective optimization and discrete multi-law supervision.

4. Verifier-Gated Multi-Candidate Framework

4.1. Architecture

Let $\theta = (R_s, L_d, L_q, \lambda_m)$ denote an electrical parameter vector and $\hat{\theta}$ its fixed controller-side value. The functions $\hat{T}_e(r) = T_e(r; \hat{\theta})$ and $\hat{v}_s(r) = v_s(r; \hat{\theta})$ use (3)–(6). The superscript eval identifies algebraic quantities evaluated with a specified test parameter vector. In electrical closed-loop simulation, $i_d(t), i_q(t)$ are integrated plant states, whereas $r_k = (i_d^\star, i_q^\star)$ is the issued reference. The vector $w = (w_T, w_V, w_I, w_C, w_S, w_M)^\mathsf{T}$ contains ranking weights; the separate coefficients $\beta_T, \ldots, \beta_E$ define the offline calibration objective. Unhatted motor coefficients in candidate formulas denote $\hat{\theta}$ unless an evaluation vector is explicitly specified. At each supervisory update instant k, the reference selector observes

$$x_k = (T_{e,k}^\star, \omega_{e,k}, V_{\mathrm{dc},k}, \hat{R}_s, \hat{L}_d, \hat{L}_q, \hat{\lambda}_m, I_{\max}, V_{\max}). \tag{13}$$

Throughout the candidate-generation, verification, and scoring equations, dependence on the stored reference r_{k-1} and candidate label m_{k-1} is implicit. The candidate generator creates a finite candidate set

$$\mathcal{A}(x_k) = \{a_1, a_2, \ldots, a_N\}. \tag{14}$$

Each candidate option returns a candidate current reference

$$r_j = (i_{d,j}^\star, i_{q,j}^\star). \tag{15}$$

The verifier constructs the verified set

$$\mathcal{A}_{\mathrm{ver}}(x_k) = \{a_j \in \mathcal{A}(x_k) : \mathcal{V}(x_k, r_j) = 1\}. \tag{16}$$

The selector then solves

$$a_k^\star = \underset{a_j \in \mathcal{A}_{\mathrm{ver}}(x_k)}{\arg\min} \; J(x_k, a_j; w), \tag{17}$$

where J is a weighted candidate score, and w is tuned offline. The selected reference is sent to the inner current controller.

Figure 1 separates the online candidate, verification, and ranking paths from the offline CEM update. The resulting verified set is a controller-model admissibility certificate under the encoded equations and margins; perturbed evaluation-model feasibility is assessed separately.

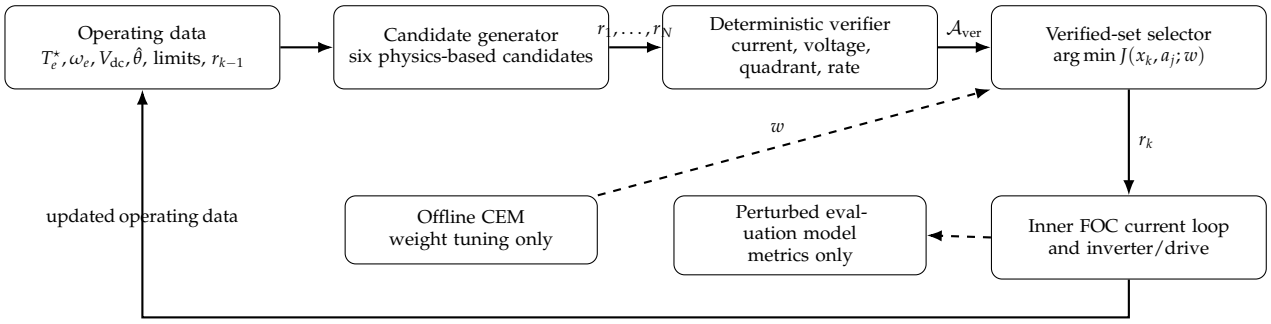


Figure 1. Proposed architecture. CEM changes only the score weights offline. The online path generates a finite candidate set, applies deterministic controller-model checks, and ranks only the verified candidates. The perturbed evaluation model is used to compute numerical metrics and is not an online verifier input. The diagram shows the nonempty verified-set path; exceptional statuses are defined in Section 4.1.

Equation (17) is used when $\mathcal{A}_{\mathrm{ver}}(x_k)$ is nonempty. The supervisory period is $T_s = 4$ ms, the initial reference is $r_{-1} = (0, 0)$ A, and the initial label is the conservative-fallback label. The first-sample switching penalty is zero. Each raw candidate is moved toward its target by at most $\Delta r_{\max} = 0.75$ A before verification.

The returned status distinguishes three outcomes. Status 0 returns a candidate passing every verifier check. If the verified set is empty, the raw conservative fallback is independently checked against the nominal electrical constraints. If it passes, status 1 returns that reference with an explicit rate-override flag; the increment guarantee is then unavailable for that update. Otherwise, status 2 returns no certified reference. A numerical placeholder for status 2 is not an actuator command. This interface exposes a request for external protection or shutdown rather than assuming that zero current is voltage-feasible. In fact, (3)–(5) give $v_s(0, 0) = |\omega_e \lambda_m|$. No override or uncertified status occurs on the four development and five primary evaluation profiles. The forced-collapse tests in Section 7.7 examine these exceptional states.

4.2. Candidate Options

The candidate set contains six deterministic candidate options evaluated on the common scalar grid. The voltage residual used by the correction option is $\widehat{\Phi}_V(i_d) = \Phi_V(i_d; \hat{\theta})$, with $i_q(i_d)$ reconstructed from the same nominal parameters. Its derivative has units V^2/A. For every exact-torque option, i_q is reconstructed from (9), and only finite motoring-branch points satisfying the current-domain test are retained. The margin-matched scalar J_ρ implementation derived from [26] is evaluated separately, so the supervisor is not benchmarked against one of its own candidates.

MTPA option. The nominal MTPA candidate is the exact-torque grid point minimizing $i_d^2 + i_q(i_d)^2$. If no current-feasible exact-torque point exists, the maximum-achievable-torque routine described below is used without a voltage restriction.

Voltage-projection option. Among exact-torque points satisfying the full nominal current and voltage limits, this option selects the point with minimum Euclidean distance from the MTPA candidate. If the exact-torque feasible set is empty, the maximum-achievable-torque routine is used with the full nominal limits.

Torque-preserving FW option. This option selects the minimum-current exact-torque point satisfying $\hat{v}_s \leq 0.90V_{\max}$. If no such point exists, the maximum-achievable-torque routine is invoked with $0.90V_{\max}$ and $\kappa_I I_{\max}$.

Voltage-minimizing option. It chooses the current-feasible exact-torque point with minimum nominal voltage magnitude. If that point exceeds $\kappa_V V_{\max}$, the maximum-achievable-torque routine is invoked with $\kappa_V V_{\max}$ and $\kappa_I I_{\max}$.

One-step residual-correction option. The MTPA candidate initializes $i_d^{(0)}$ in

$$i_d^{(1)} = i_d^{(0)} - \alpha \frac{\widehat{\Phi}_V(i_d^{(0)})}{\widehat{\Phi}'_V(i_d^{(0)}) + \mathrm{sgn}_0\Big(\widehat{\Phi}'_V(i_d^{(0)})\Big)\epsilon}, \tag{18}$$

where $\mathrm{sgn}_0(z) = 1$ for $z \geq 0$ and -1 otherwise. The voltage-squared residual derivative is approximated by a centered grid difference when both neighboring points are available and by the available one-sided difference at a bound. The implementation uses $\alpha = 0.80$ and $\epsilon = 10^{-8}$ V^2/A. The correction is applied only when $\widehat{\Phi}_V(i_d^{(0)}) > 0$ and the magnitude of the regularized denominator exceeds ϵ; otherwise $i_d^{(1)} = i_d^{(0)}$. The resulting $i_d^{(1)}$ is clipped to $[-I_{\max}, 0]$, and $i_q^{(1)}$ is recomputed from (9).

Conservative fallback option. At each grid point, the nonnegative i_q value is limited by the requested torque, the current circle, and the positive root of the full voltage equation. The selected point minimizes the squared normalized torque shortfall plus the regularization $10^{-5}(i_d^2 + i_q^2) + 10^{-7}v_s^2$, subject to $\|r\|_2 \leq 0.90I_{\max}$ and $v_s \leq 0.90V_{\max}$. This option prioritizes an achievable derated torque and deterministic behavior over exact torque tracking.

For all six options, the raw target r_j^{raw} is rate-limited before verification according to

$$r_j = r_{k-1} + \min\left(1, \frac{\Delta r_{\max}}{\|r_j^{\mathrm{raw}} - r_{k-1}\|_2}\right)(r_j^{\mathrm{raw}} - r_{k-1}),$$

with the unchanged target used when the denominator is zero.

4.3. Verifier

The verifier is deterministic. For a candidate r_j, it checks the current, voltage, sign, and rate constraints

$$\|r_j\|_2 \leq \kappa_I I_{\max}, \tag{19}$$

$$\hat{v}_s(x_k, r_j) \leq \kappa_V V_{\max}, \tag{20}$$

$$i_{q,j}^{\star} \geq 0, \tag{21}$$

$$\|r_j - r_{k-1}\|_2 \leq \Delta r_{\max}, \tag{22}$$

where $\kappa_I = 0.985$, $\kappa_V = 0.935$, and $\Delta r_{\max} = 0.75$ A per 4 ms update.

Here, $\hat{v}_s$ is computed with the nominal controller-side parameters. The verifier also checks $i_d \leq 0$ and $\lambda_m + (L_d - L_q)i_d \geq 10^{-8}$ Wb in addition to the explicit $i_q \geq 0$ test in (21). The proposition therefore establishes ranking-independent controller-model admissibility. Scenario-specific evaluation parameters are used only after selection to test how that model-side certificate transfers under dc-link and parameter perturbations.

Proposition 1 (Ranking-weight-independent verifier admissibility)**.** *If $\mathcal{A}_{\mathrm{ver}}(x_k)$ is nonempty, then for any weight vector w—and, more generally, for any ranking or tie-breaking rule whose output belongs to $\mathcal{A}_{\mathrm{ver}}(x_k)$—the executed reference satisfies every current, voltage-margin, branch, denominator, and increment check encoded by $\mathcal{V}$ at sample k.*

Proof. Membership in $\mathcal{A}_{\mathrm{ver}}(x_k)$ is equivalent to $\mathcal{V}(x_k, r_j) = 1$. The ranking rule changes only which member of this nonempty set is selected and therefore cannot invalidate any encoded check. □

Proposition 2. *If every executed reference, including any empty-set fallback, satisfies (22), then the executed reference sequence is incrementally bounded by $\Delta r_{\max}$.*

Proof. Applying (22) to every executed sample gives $\|r_k - r_{k-1}\|_2 \le \Delta r_{\max}$ for all k. A status-1 emergency records violation of this premise. The five primary profiles have no override; the abrupt-collapse diagnostic deliberately includes overrides. □

Let $\mathcal{C}_k$ be the nominal electrical set defined by (19)–(21), $i_d \le 0$, and the positive denominator guard, excluding (22). At fixed operating data, it is convex: the current norm is convex, the voltage vector is affine in r, and the branch conditions are linear. Suppose the raw fallback r_F belongs to $\mathcal{C}_k$. Write $r(\alpha) = r_{k-1} + \alpha(r_F - r_{k-1})$, and let $\alpha_{\max} \in [0, 1]$ give the largest step allowed by (22). If $r(\alpha_{\max}) \notin \mathcal{C}_k$, no smaller α on this segment is feasible. Otherwise, a feasible smaller point and r_F would, by convexity, make $r(\alpha_{\max})$ feasible, a contradiction. Consequently, scanning 101 points backward on this same segment cannot recover a rate-feasible point under the stated premise; the status-1 branch instead returns the electrically rechecked fallback directly.

This segment result is not a global infeasibility test. With $\mathbb{B}(r_{k-1}, \Delta r_{\max})$ denoting the closed Euclidean increment ball, an electrically and increment-feasible reference exists exactly when $\mathcal{C}_k \cap \mathbb{B}(r_{k-1}, \Delta r_{\max}) \neq \varnothing$. If $\mathcal{C}_k$ is nonempty, this is equivalent to $\mathrm{dist}(r_{k-1}, \mathcal{C}_k) \le \Delta r_{\max}$. Status 2 reports failure to return a certified candidate, not a proof that this intersection is empty.

Proposition 3 (Fixed-reference box verification)**.** *Fix r, ω_e, and $V_{\max}$. For an independent parameter box $\Theta = \prod_{j=1}^{4} [\theta_j^-, \theta_j^+]$, let $\mathrm{vert}(\Theta)$ be its 16 vertices, and let $v(r;\theta) = (v_d, v_q)^{\mathsf{T}}$ be the quasi-steady voltage vector. Then,*

$$\max_{\theta \in \Theta} \|v(r;\theta)\|_2 = \max_{\theta \in \mathrm{vert}(\Theta)} \|v(r;\theta)\|_2.$$

Proof. Equations (3) and (4) are affine in θ for fixed r and ω_e. Each point of Θ is a convex combination of its vertices, so its voltage vector is the same convex combination of vertex voltage vectors. Convexity of the Euclidean norm bounds that value by the largest vertex value. The reverse inequality follows because every vertex belongs to Θ. □

Proposition 3 supports the fixed-reference uncertainty evaluation in Section 7.5. It neither changes the online nominal gate nor certifies closed-loop trajectories by testing only 16 dynamic plants. In particular, (20) establishes a voltage condition on the selected reference under $\hat{\theta}$; plant-current tracking error, derivative voltages, and inverter effects belong to the separate dynamic evaluation.

4.4. Candidate Score

Among verified candidates, the selector minimizes

$$J(x_k, a_j; w) = w_T E_T(r_j) + w_V E_V(r_j) + w_I E_I(r_j) + w_C E_C(r_j) + w_S E_S(r_j) + w_M E_M(r_j), \tag{23}$$

where

$$E_T = \left(\frac{\hat{T}_e(r_j) - T^\star_{e,k}}{\max(1, |T^\star_{e,k}|)} \right)^2, \tag{24}$$

$$E_V = \left(\frac{\hat{v}_s(x_k, r_j)}{V_{\max}} \right)^2, \tag{25}$$

$$E_I = \left(\frac{\|r_j\|_2}{I_{\max}} \right)^2, \tag{26}$$

$$E_C = \left(\frac{\|r_j - r_{k-1}\|_2}{0.2 I_{\max}} \right)^2, \tag{27}$$

$$E_S = \mathbf{1}[m_j \neq m_{k-1}], \tag{28}$$

with m_j denoting the candidate label. The margin term is

$$E_M = \left[\max\left(0, \frac{\hat{v}_s(x_k, r_j)/V_{\max} - 0.90}{0.10}\right) \right]^2. \tag{29}$$

The denominator in (24) is 1 N m when $|T^\star_{e,k}| < 1$ N m. Equations (25) and (29) use the same nominal controller-side voltage as the verifier; the perturbed evaluation-model voltage is not available to the online score. The two voltage terms have different roles: E_V penalizes overall utilization, whereas E_M activates above 0.90 per unit. At the first sample, $E_S = 0$ for every option. Thereafter, the previous executed label is used. All six weights are strictly positive because CEM samples them through the exponential map in (30). The score controls ranking inside the verified set and cannot override (19)–(22). Candidate-score ties are resolved by the first candidate in the fixed order listed in Section 4.2.

5. CEM-Based Weight Tuning

The cross-entropy method is used only for offline tuning of w [27,28]. It does not modify the candidate generator, the verifier, or the inner current controller.

At CEM iteration n, candidate weight vectors are sampled from a log-normal distribution:

$$w_n^{(m)} = \exp(z_n^{(m)}), \qquad z_n^{(m)} \sim \mathcal{N}(\mu_n, \Sigma_n), \qquad m = 1, \ldots, M. \tag{30}$$

Each vector is evaluated on the development profiles by

$$\mathcal{J}(w) = \frac{1}{N_D K} \sum_{\ell=1}^{N_D} \sum_{k=1}^{K} \left(\beta_T e_T^2 + \beta_V \psi_V + \beta_I i_s^2 + \beta_C \Delta r^2 + \beta_S \Delta m + \beta_E e_\varnothing \right), \tag{31}$$

where $N_D = 4$, $e_T = (T_e^{\text{eval}} - T_e^\star)/\max(1, |T_e^\star|)$, $u = v_s^{\text{eval}}/V_{\max}$, $\psi_V = \max(0, u - 0.90)^2 + 50\max(0, u - 1)^2$, $i_s^2 = \|r_k\|_2^2/I_{\max}^2$, $\Delta r^2 = \|r_k - r_{k-1}\|_2^2/\Delta r_{\max}^2$, $\Delta m = \mathbf{1}[m_k \neq m_{k-1}]$ with zero first-sample value, and $e_\varnothing$ indicates an empty verified set. The coefficients are $(\beta_T, \beta_V, \beta_I, \beta_C, \beta_S, \beta_E) = (1, 25, 0.03, 0.03, 0.01, 2)$.

The CEM population is $M = 20$, the elite count is four, and eight fixed iterations are used. The initial distribution is $\mu_0 = \log([12, 1, 2, 0.8, 0.3, 4])$ and $\Sigma_0 = 0.70^2 I_6$. At each

iteration, the elite mean and covariance are blended with the preceding values using a smoothing coefficient of 0.70; $0.05^2 I_6$ is added as a covariance floor. The random seed is 20260731. The lowest-cost sampled vector observed during the eight iterations is frozen; there is no test-profile-based early stopping or retuning. The resulting vector is

$$w_{\mathrm{CEM}} = [16.600,\ 1.254,\ 1.366,\ 0.969,\ 0.101,\ 7.196], \tag{32}$$

corresponding to the torque, voltage, current, change, switching, and margin terms. The manual vector is $[12, 1, 2, 0.8, 0.3, 4]$. On the four development profiles, CEM reduced the mean rollout cost from 0.033396 to 0.030964, a 7.28% reduction.

Four deterministic development profiles were used for CEM and voltage-feedback-gain tuning; five different numerical profiles were frozen for final evaluation. The profiles use related operating categories but different commands, transition times, voltages, and mismatch factors. Accordingly, the final tables are held out from numerical tuning, but they do not establish broad statistical or structural generalization. The primary profiles and the ranking-weight panel are deterministic descriptive evaluations. Across-profile standard deviations quantify heterogeneity among these chosen profiles, not uncertainty for a sampled population.

6. Numerical Simulation Setup

6.1. Machine and Inverter Parameters

The nominal machine values are taken from the published standard IPMSM case in [23]; the same parameter set is also used in [24,26]. That source reports four poles, $R_s = 2.48\ \Omega$, $L_d = 74.98$ mH, $L_q = 113.91$ mH, $\lambda_m = 0.193$ Wb, $I_{\max} = 10$ A, and $V_{\max} = 196.7$ V. Under the $V_{\max} = V_{\mathrm{dc}}/\sqrt{3}$ convention used here, the corresponding nominal dc-link voltage is $V_{\mathrm{dc}} = 340.694$ V. The source paper uses $p = 4$ poles; the equivalent value in (6) is two pole pairs. Table 2 lists the parameters and numerical settings.

The algebraic benchmark updates the reference at 4 ms intervals. The separate electrical closed-loop model in Section 6.6 integrates current dynamics with a faster current loop and switching inverter. The 4 ms value remains a chosen supervisory period, not a measured target-controller execution time.

Table 2. Nominal IPMSM, inverter, verifier, and numerical settings used in the quasi-steady study.

Parameter	Symbol	Value
Stator resistance	R_s	2.48 Ω
Direct-axis inductance	L_d	74.98 mH
Quadrature-axis inductance	L_q	113.91 mH
PM flux linkage	λ_m	0.193 Wb
Pole pairs (four-pole source model)	p	2
Current limit	$I_{\max}$	10 A
Nominal dc-link voltage	V_{dc}	340.694 V
Nominal voltage-vector limit	$V_{\max}$	196.7 V $= V_{\mathrm{dc}}/\sqrt{3}$
Verifier current/voltage margins	κ_I, κ_V	0.985, 0.935
Supervisory period/samples	T_s, K	4 ms, 251
Reference-increment limit	$\Delta r_{\max}$	0.75 A/update
Scalar grid	i_d	$[-10, 0]$ A, 801 points
Denominator guard	—	10^{-8} Wb
Scalar-comparator transition speed	ω_a	206.90 rad/s [23]

6.2. Compared Methods

Six primary methods were compared using the same nominal motor model, scalar grid, $r_{-1} = (0, 0)$ A, and 0.75 A/update external reference-increment cap.

1. **MTPA + rule-based FW:** MTPA is used while it satisfies the full nominal voltage limit. Otherwise, the method selects the exact-torque grid point nearest to MTPA that satisfies the resistance-neglected voltage equation and the current limit. If that set is empty, it invokes the full-limit maximum-achievable-torque routine.
2. **Voltage-feedback FW:** The nominal voltage-utilization error is $e_v = \hat{v}_s / V_{\max} - 0.985$. The integral state is updated with $T_s e_v$ and clamped to $[-0.10, 0.25]$. The weakening amount is $\max(0, K_P e_v + K_I \max(0, \xi))$, with $K_P = 8$ A/pu and $K_I = 5$ A/(pu s); i_d is shifted negatively from MTPA, clipped to $[-I_{\max}, 0]$, and i_q is reconstructed and current-limited. The gains were selected from $K_P \in \{2, 4, 6, 8\}$ and $K_I \in \{5, 15, 30\}$ on the development profiles.
3. **Projection reference:** This is the voltage-projection candidate defined in Section 4.2 without verifier margins or score-based selection.
4. **Margin-matched scalar J_ρ comparator:** Following the closest antecedent [26], the exact-torque scalar points satisfying the effective limits $\kappa_I I_{\max}$ and $\kappa_V V_{\max}$ are evaluated using
$$J_\rho = (1 - \rho) \frac{i_d^2 + i_q^2}{I_{\max}^2} + \rho \frac{v_d^2 + v_q^2}{V_{\max}^2}, \qquad \rho = \frac{\omega_e^2}{\omega_e^2 + \omega_a^2},$$
with $\omega_a = 206.90$ rad/s. The minimum-J_ρ point is selected; if the exact-torque feasible interval is empty, the same maximum-achievable-torque routine is used. The common external increment cap is then applied. The comparator reproduces the antecedent's feasible-set and fixed-objective decision logic on the common 801-point grid; it is not a timing or numerical-accuracy reproduction of that paper's safeguarded Newton/bisection implementation.
5. **Verifier multi-candidate:** The six candidates and verifier are used with manual weights $[12, 1, 2, 0.8, 0.3, 4]$.
6. **Verifier multi-candidate + CEM:** The same candidates and verifier are used with the frozen weights in (32).

A separate **score-only multi-candidate + CEM** ablation uses the same six rate-limited candidates and the same frozen CEM weights but removes all verifier masking. It is not included as a separate primary method because, for the frozen weights and five frozen evaluation profiles, it selected the same trajectory as the gated version; Section 7.3 reports this result directly.

The labels denote these implementations only. In particular, the scalar comparator is one registered mixed objective from the more general framework in [26], and the voltage-feedback baseline is not intended to represent every anti-windup or measured-voltage implementation in that literature.

6.3. Development and Evaluation Profiles

Let $\tau = t/1\ \text{s} \in [0, 1]$, and define the cubic transition $s_{a,b}(\tau) = z^2(3 - 2z)$ with $z = \text{clip}[(\tau - a)/(b - a), 0, 1]$. Each profile contains 251 samples from $t = 0$ to 1 s. The four development profiles were as follows:

- $T_e^\star = 6.2 + 0.6 \sin(2\pi\tau)$, $\omega_e = 135 + 150 s_{0.08,0.92}$, nominal dc link and parameters;
- $T_e^\star = 3.2 + 2.8(s_{0.25,0.32} - s_{0.70,0.77})$, $\omega_e = 285$ rad/s, nominal dc link and parameters;
- $T_e^\star = 6.8$ N m, $\omega_e = 240$ rad/s, $V_{\text{dc}} = V_n - (V_n - 292)(s_{0.28,0.35} - s_{0.66,0.73})$, and evaluation multipliers $(R_s, L_d, L_q, \lambda_m) = (1.08, 0.97, 1.03, 0.97)$; and

- $T_e^\star = 5 + 1.6\sin(4\pi\tau + 0.3) + 0.7\sin(10\pi\tau)$, $\omega_e = 205 + 70\sin(2\pi\tau - 0.4) + 30\sin(6\pi\tau)$, a sag to 300 V between the $s_{0.42,0.48}$ and $s_{0.62,0.68}$ transitions, and multipliers $(1.10, 0.97, 1.03, 0.97)$.

Here, $V_n = 340.694$ V. Torque commands are clipped to $[0.2, 9.2]$ N m and speeds to $[110, 330]$ rad/s in all profiles.

Table 3 specifies the five frozen evaluation profiles and the stressor examined by each. The base-speed transition increases electrical speed from 150 to 305 rad/s at a fixed 7.0 N m command, testing the shift into voltage-limited operation. The high-speed torque step instead changes the command from 3.5 to 6.5 N m and back at 300 rad/s, testing a change in torque demand without a speed change. The dc-link-sag case reduces the available voltage at fixed torque and speed. These three cases retain nominal machine parameters. The mismatch case evaluates the references with perturbed electrical parameters while leaving the controller model nominal, and the combined cycle applies simultaneous torque, speed, dc-link, and parameter variations. Thus, the panel distinguishes individual constraint stresses from their combined effect, using the same frozen controller settings throughout.

Table 3. Five evaluation profiles frozen after tuning. The mismatch multipliers are applied only when computing evaluation-model torque and voltage; all reference generators use nominal parameters.

Profile	$T_e^\star$ (N m)	ω_e (rad/s)	V_{dc} (V)	Evaluation Multipliers $(R_s, L_d, L_q, \lambda_m)$
Base-speed transition	7.0	$150 + 155 s_{0.05,0.95}$	V_n	$(1, 1, 1, 1)$
High-speed torque step	$3.5 + 3.0(s_{0.30,0.36} - s_{0.68,0.74})$	300	V_n	$(1, 1, 1, 1)$
Dc-link sag	7.4	250	$V_n - (V_n - 278)(s_{0.30,0.36} - s_{0.64,0.70})$	$(1, 1, 1, 1)$
Parameter mismatch	$6.2 + 0.4\sin(2\pi\tau)$	$225 + 70 s_{0.10,0.90}$	V_n	$(1.20, 0.95, 1.05, 0.95)$
Combined drive cycle	$5.2 + 1.8\sin(4\pi\tau) + 0.8\sin(10\pi\tau + 0.5)$	$210 + 78\sin(2\pi\tau - 0.6) + 35\sin(6\pi\tau + 0.2)$	Sag to 292 using $s_{0.43,0.49} - s_{0.64,0.70}$	$(1.10, 0.97, 1.03, 0.97)$

The parameter-mismatch multipliers follow the combined $R_s + 20\%$, $\lambda_m - 5\%$, $L_d - 5\%$, and $L_q + 5\%$ sensitivity case reported for the same benchmark motor in [24].

6.4. Metrics, Ablations, and Reproducibility Checks

Let $v_s^{\mathrm{eval}}(k)$ and $T_e^{\mathrm{eval}}(k)$ denote quantities computed with the scenario-specific evaluation parameters. The primary metrics are torque root mean square error (RMSE) (N m); voltage-violation rate $\mathrm{VR} = 100K^{-1}\sum_k \mathbf{1}[v_s^{\mathrm{eval}}(k) > V_{\max}(k)]$; current-violation rate $\mathrm{IR} = 100K^{-1}\sum_k \mathbf{1}[\|r_k\|_2 > I_{\max}]$; average voltage-margin index $\mathrm{VMI} = K^{-1}\sum_k [1 - v_s^{\mathrm{eval}}(k)/V_{\max}(k)]$; mean-squared-current index $\mathrm{MSI} = K^{-1}\sum_k \|r_k\|_2^2$; and mode-switching ratio $\mathrm{MSR} = 100(K-1)^{-1}\sum_{k=2}^{K} \mathbf{1}[m_k \neq m_{k-1}]$. The maximum inter-update reference increment is also recorded. For a gated rollout, $\mathrm{RR} = 100(6K)^{-1}\sum_k n_{\mathrm{rej}}(k)$, where n_{rej} is the number of rejected candidates among the six.

Smaller RMSE, VR, IR, MSI, and MSR are preferable, whereas larger VMI indicates greater average voltage headroom. VMI is an average rather than a minimum-margin metric. MSI is a current-use surrogate and is not measured drive efficiency.

Two diagnostic ablations are used. First, the frozen CEM score is evaluated with and without verifier masking while all candidate values and rate limits are held fixed. Second, ranking-weight dependence is stressed with 128 fixed-seed vectors $w^{(q)} = w_{\mathrm{CEM}} \odot 10^{u^{(q)}}$, $u_\ell^{(q)} \sim \mathcal{U}[-2, 2]$, using seed 20260912. Each positive score weight is therefore perturbed independently over four decades. This panel is a deliberately broad diagnostic stress test, not a probability model for deployed weights. For each vector, the score-only and gated selectors are run on the same five evaluation profiles. The reported quantities include the fraction of weight vectors producing any controller-model current, voltage, or increment violation and the mean controller- and evaluation-model voltage-violation rates.

6.5. Diagnostic and Uncertainty Protocols

The candidate ablation compares the full set with each of its six leave-one-out variants using the same frozen weights, margins, profiles, and tie rule. Removing label 6 removes it from ordinary ranking only; the emergency handler remains available. The outputs include changed-reference counts, performance metrics, selection counts, and sole-verified-candidate counts. Grid refinement uses 201, 401, 801, 1601, 3201, and 6401 points with no weight retuning. The 6401-point result is a finer-grid comparison, not an exact continuous optimizer.

For uncertainty analysis, Θ_δ contains independent multipliers in $[1-\delta, 1+\delta]$ around each nominal electrical parameter. All 16 vertices are evaluated for $\delta \in \{0.05, 0.10, 0.20\}$ at each of the 1255 frozen reference samples. These boxes are synthetic diagnostic envelopes, not measured production tolerances or calibrated temperature ranges. One-at-a-time multipliers are $0.7, 0.8, \ldots, 1.3$. Controller parameters, reference trajectories, and score weights remain unchanged; the specified box replaces, rather than compounds, the profile's evaluation multipliers. Worst utilization and violation duration are descriptive results over this finite panel. The same vertex evaluation is also applied to every ablated reference trajectory.

To distinguish unavailable torque from controller derating, $T^{\mathrm{eval}}_{\mathrm{cap},G}(k)$ denotes the largest evaluation-model torque on a 6401-point i_d grid under the full current and voltage limits. At each i_d, the largest nonnegative i_q allowed by the current circle and voltage quadratic is used, without the fallback regularization. This is a grid-estimated capability, not a continuous optimum certificate. The capability-capped shortfall is

$$\max\Big(0, \min\{T^{\star}_{e,k}, T^{\mathrm{eval}}_{\mathrm{cap},G}(k)\} - T^{\mathrm{eval}}_{e}(r_k)\Big).$$

All cases, including failures, are retained in the analysis. The numerical settings and evaluation procedures are specified above. No preregistration or field-reliability inference is associated with these diagnostic tests.

6.6. Sampled Electrical Closed-Loop Model

The electrical plant integrates (1) and (2), retaining $L_d\dot{i}_d$ and $L_q\dot{i}_q$. Rotor electrical speed is prescribed by the profile, so the simulation concerns electrical current regulation under imposed speed rather than mechanical speed tracking. The nominal controller and the evaluation plant use the parameter conventions defined in Section 4.1. The supervisory update is 4 ms; both the current-control period T_c and PWM period are 100 µs. The initial plant current and the two PI integral states are zero.

Let n_{c} be the current-loop sample index, $i_j^{\mathrm{meas}}[n_{\mathrm{c}}]$ the one-period-delayed current measurement, and $e_j[n_{\mathrm{c}}] = i_j^{\star}[n_{\mathrm{c}}] - i_j^{\mathrm{meas}}[n_{\mathrm{c}}]$ for $j \in \{d, q\}$. Independent zero-mean Gaussian noise with standard deviation $\sigma_i = 0.02$ A is included only in the nonideal mode. The nominal tuning bandwidth is $\omega_c = 2\pi(200)$ rad/s, with $K_{P,d} = \hat{L}_d\omega_c$, $K_{P,q} = \hat{L}_q\omega_c$, and $K_{I,d} = K_{I,q} = \hat{R}_s\omega_c$. The PI integral output ξ_j has units of volts. The commanded voltage vector is

$$v^{\mathrm{cmd}} = \begin{bmatrix} K_{P,d}e_d + \xi_d - \omega_e\hat{L}_q i_q^{\mathrm{meas}} \\ K_{P,q}e_q + \xi_q + \omega_e(\hat{L}_d i_d^{\mathrm{meas}} + \hat{\lambda}_m) \end{bmatrix}.$$

Radial saturation gives v^{lim} with $\|v^{\mathrm{lim}}\|_2 \le V_{\max}$; a zero commanded vector is left unchanged. Back-calculation uses

$$\xi_j[n_{\mathrm{c}}+1] = \xi_j[n_{\mathrm{c}}] + T_c\{K_{I,j}e_j[n_{\mathrm{c}}] + \omega_c(v_j^{\mathrm{lim}}[n_{\mathrm{c}}] - v_j^{\mathrm{cmd}}[n_{\mathrm{c}}])\}.$$

The gains have units V/A and V/(A s), respectively. This PI law and its settings are common to every compared reference generator.

Inverse Park transformation and min–max common-mode injection produce three duty ratios, held for one PWM period and applied with one further period of actuation delay. A centered triangular carrier drives a two-level inverter. In the nonideal switching mode, a 2 µs complementary-device blanking interval is modeled at transitions; during blanking, the conducting diode is selected by phase-current sign. A 0.8 V sign-dependent device drop is also included. The current-loop bandwidth, dead time, device drop, and noise level are explicit simulation assumptions, not measurements of the benchmark motor or inverter.

The three modes are averaged ideal inverter, ideal switching inverter, and nonideal switching inverter. Each mode is evaluated for projection, scalar, and supervisor references on the five primary profiles plus a hard torque step: 3.5 to 6.5 N m at 0.30 s and back at 0.70 s, with $\omega_e = 300$ rad/s and nominal dc link. This gives 54 primary runs. The continuous optimizer defined in Section 6.8 is additionally evaluated in the nonideal switching mode on all six profiles, giving 60 electrical closed-loop runs in total. A single fixed noise realization per profile is shared by methods; seeds are $20261017 + j$ for profile indices $j = 0, \ldots, 5$. These are not independent repeated-noise trials.

Fourth-order Runge–Kutta integration uses a 1 µs step with rotor-angle-dependent transformations at its internal stages. Operating inputs are linearly interpolated between 100 µs samples. Current peaks and violations are recorded at integration steps; torque RMSE is computed at current-loop samples over 0–1 s, including startup. The nonideal sag cases are repeated at 0.5 µs using the same noise. An independent constant-parameter matrix-exponential calculation checks the electrical right-hand side and integration for three prescribed voltage/speed cases.

Voltage demand is $u_{\mathrm{dem}} = \|v^{\mathrm{cmd}}\|_2 / V_{\max}$, and saturation duration is the percentage of current-loop periods with $u_{\mathrm{dem}} > 1$. The limited command, PWM-period-average applied voltage, and quasi-steady voltage demand are stored separately. Instantaneous switching vectors are not classified against the fundamental-voltage limit in (8). This prevents controller saturation from being mislabeled as an applied-voltage constraint violation.

6.7. Online Cost and Emergency Tests

The single-update timing path includes candidate generation, rate limiting, verification, ranking, and any emergency decision. Compilation, warm-up, plotting, file output, and offline CEM are excluded. The supervisor and scalar comparator are timed on the same 1255 operating-input/history tuples for ten repetitions, rotating execution order; 4000 further supervisor calls cover four abrupt-collapse inputs. Wall-clock samples are retained without outlier removal. The environment is an Intel Xeon Platinum 8573C host, Linux x86-64, Python 3.13.5, NumPy 2.3.5, and Numba 0.65.1 with fast-math disabled. The process is pinned to one logical CPU; this does not isolate it from all host scheduling effects.

The implementation uses double precision. The principal stored numeric buffers are inventoried separately from process peak resident memory. Neither a host maximum nor a buffer subtotal is a target-processor WCET or total embedded RAM requirement. The collapse diagnostic fixes $T_e^\star = 7.4$ N m and $\omega_e = 250$ rad/s and tests dc-link values of 150, 100, 50, and 5 V from a common preceding reference.

The additional numerical analyses were implemented in Python/Numba using the algebraic reference-generation laws and the sampled electrical model specified in Sections 4 and 6.6. The reported results are computed outputs obtained under the numerical settings and evaluation protocols described above.

6.8. Equal-Budget Calibration and Continuous Optimization

A separate comparison evaluates tunable methods using one common external development criterion. For a reference r at update k, define

$$\ell_c(r;k,\theta) = \left[\frac{T_e(r;\theta) - T^\star_{e,k}}{\max(1\,\mathrm{N\,m}, |T^\star_{e,k}|)}\right]^2 + 25\{\max(u-0.90,0)^2 + 50\max(u-1,0)^2\} + 0.03\frac{\|r\|_2^2}{I_{\max}^2} + 0.03\frac{\|r-r_{k-1}\|_2^2}{\Delta r_{\max}^2},$$

where $u = v_s(r;\theta)/V_{\max,k}$. The development cost $\mathcal{J}_c$ is the mean of this loss over all samples and the four development profiles, using each profile's evaluation parameters. Its coefficients are fixed across methods. Unlike (31), it contains neither an option-label penalty nor a candidate-empty-set penalty, which would not have the same interpretation in all architectures. Nonfinite rollouts receive a cost of 10^6; all evaluations and selection records are retained.

The supervisor, a specified voltage-feedback implementation, and the scalar schedule each receive 20 sampled settings per iteration for eight CEM iterations, plus one initial setting: 161 development-cost calls per method. Each uses four elites, 0.70 smoothing, initial log-coordinate standard deviation 0.70, a $0.05^2 I$ covariance floor, and seed 2026101701. Bounds are enforced in log coordinates. The supervisor fixes the torque weight to one and tunes the five positive weight ratios in $[10^{-4}, 100]$, initialized from the manual weights divided by 12. The scalar tunes $\omega_a \in [50, 600]$ rad/s from 206.90 rad/s. The five evaluation profiles are not used for setting selection.

The matched-margin feedback variant is defined explicitly rather than identified with every voltage-feedback method. Its error is the nominal voltage at the previous issued reference, evaluated at current operating data, divided by $V_{\max,k}$, minus κ_V. Its integral state starts at zero and is clamped to $[-0.10, 0.25]$ pu s. It shifts the MTPA i_d by $\max(0, K_P e_v + K_I \max(0,\xi))$, reconstructs i_q from (9), applies the $\kappa_I I_{\max}$ circle and the common increment cap, and receives no verifier masking. The tuning ranges are $K_P \in [0.1, 60]$ A/pu and $K_I \in [0.1, 300]$ A/(pu s), initialized at 8 and 5. This precisely defined comparison variant is distinct from the primary full-limit feedback setting in Section 6.2; it uses no evaluation-model information online.

A direct two-variable nonlinear-programming (NLP) comparator minimizes $\ell_c(r;k,\hat{\theta})$ at each update subject to the same nominal electrical and reference-increment constraints. It does not select from the six-option library. The decision vector is $\zeta = r/I_{\max}$, with bounds $\zeta_d \in [-\kappa_I, 0]$ and $\zeta_q \in [0, \kappa_I]$. SciPy 1.17.0 SLSQP uses analytic objective gradients and constraint Jacobians, a function tolerance of 10^{-9}, and at most 100 iterations per start [29]. The two starts are the previous reference and its rate-limited move toward the conservative fallback. A 10^{-7} normalized backoff in the current, voltage, and increment bounds reduces boundary-rounding failures. Every finite returned point is independently tested by the unmodified nominal verifier, regardless of the solver's success flag; the feasible point with lower objective is selected. If both starts fail this check, a verified previous reference is retained, otherwise a checked electrical fallback is returned with an override, or no certified reference is returned. This is a local nonconvex search, not a global optimality certificate. The criterion is fixed and has no additional fitted weights.

Contemporary MTPA/MTPV analysis can incorporate experimentally calibrated nonlinear flux-linkage maps, as in Fan et al. [30]. Those maps encode saturation and cross-coupling absent from this constant-parameter benchmark. The NLP comparison therefore tests continuous optimization under identical available equations and limits; it is not a numerical reproduction of that nonlinear-map study. This preserves a like-for-like quantitative comparison while identifying the physical modeling scope that remains different.

7. Results

7.1. Aggregate Comparison

Table 4 reports the unweighted mean of each metric across the five frozen evaluation profiles. The three conventional non-verifier baselines produced average voltage-violation rates of 10.2–38.3%. The margin-matched scalar comparator and both verifier-gated selectors reduced evaluation-model VR and IR to 0.0% across all five profiles. All six primary methods respected the common 0.75 A increment cap.

Table 4. Average performance across the five evaluation profiles. MSI is in A^2. RR applies only to the proposed gated selectors. Values are deterministic descriptive results.

Method	RMSE (N m)	VR (%)	IR (%)	VMI	MSI (A^2)	MSR (%)	RR (%)
MTPA + rule FW	0.794	34.7	0.0	0.117	49.70	0.32	—
Voltage-feedback FW	0.745	38.3	0.0	0.105	49.66	0.00	—
Projection reference	0.797	10.2	0.0	0.128	50.60	0.00	—
Margin-matched scalar J_ρ	0.938	0.0	0.0	0.201	53.33	0.00	—
Verifier multi-candidate	0.979	0.0	0.0	0.175	50.89	0.56	14.24
Verifier multi-candidate + CEM	0.979	0.0	0.0	0.182	51.49	0.64	14.24

Relative to the closest margin-matched scalar comparator, the CEM-tuned supervisor exchanged 4.4% higher torque RMSE and a 0.019 lower VMI for 3.46% lower MSI, while both methods maintained zero VR and IR. This result quantifies the design trade-off rather than collapsing it into a single aggregate winner: the scalar schedule favored torque tracking and voltage headroom, whereas multi-candidate selection used less squared current.

Relative to the manual score, CEM moved the held-out operating point toward greater voltage headroom: VMI increased from 0.175 to 0.182, VR and IR remained zero, and RMSE changed by less than 3×10^{-4} N m. MSI increased from 50.89 to 51.49 A^2 and MSR from 0.56% to 0.64%. Thus, the 7.28% reduction in the fixed development objective produced the intended multi-term calibration but not a monotonic improvement in every separately reported metric.

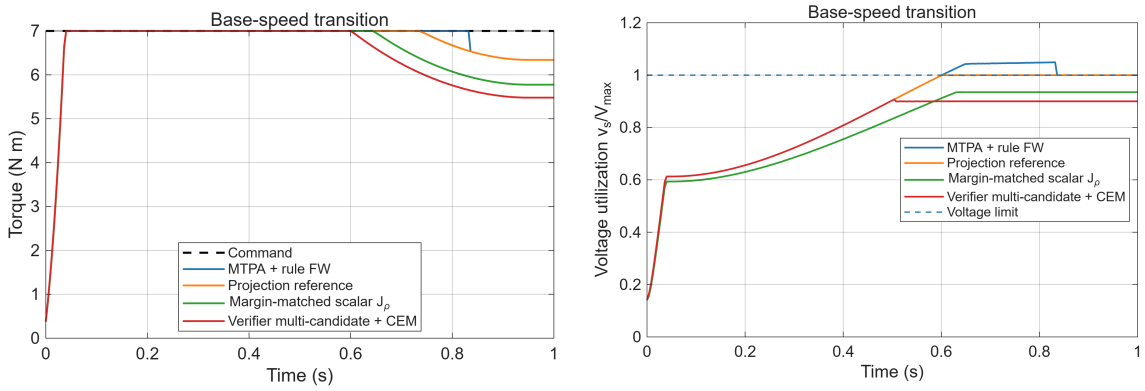
Across the five chosen profiles, the scalar comparator has standard deviations of 0.339 N m in RMSE, 11.43 A^2 in MSI, and 0.103 in VMI. The corresponding supervisor values are 0.361 N m, 12.47 A^2, and 0.083. These sample standard deviations use the profile as the unit of analysis ($n = 5$, divisor $n - 1$) and describe variation among the selected evaluation profiles rather than confidence intervals for field operation.

7.2. Profile-Level Results

Table 5 reports four representative profiles for the rule-based FW, projection, margin-matched scalar, and CEM-tuned supervisory methods. The scalar comparator and supervisor use the same current and voltage margins and the same external reference-increment cap, so their differences reflect the selected decision laws rather than different effective limits. Figures 2–4 complement these profile-level metrics with torque and voltage-utilization traces for three of the four tabulated profiles. The plotted lines connect quasi-steady samples; they do not represent simulated inner-loop or switching dynamics.

Figure 2 shows the base-speed transition at the constant 7.0 N m command specified in Table 3. After the common rate-limited initialization, the torque traces in Figure 2a separate as increasing speed makes voltage limitation active. Figure 2b shows the rule-based method exceeding unit voltage utilization, while the projection method follows the full nominal boundary and the two margin-based methods remain below it. Table 5 quantifies this behavior: rule-based FW has 23.1% VR, whereas the other three displayed methods have zero VR. Relative to the scalar comparator, the supervisor reduces MSI from 61.81 to 59.57 A^2, with RMSE increasing from 1.042 to 1.140 N m and average VMI decreasing from 0.212 to 0.205. The lower late-profile voltage utilization of the supervisor should

therefore not be confused with a larger profile-average VMI. Together, the two panels and the table identify a lower-current operating choice that preserves voltage admissibility while accepting additional torque error.



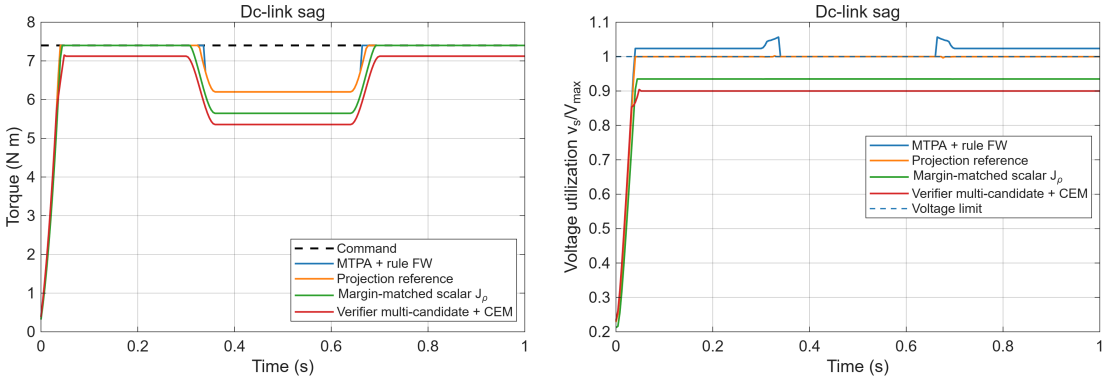
(**a**) Torque (**b**) Voltage utilization

Figure 2. Base-speed transition at $T_e^\star = 7.0$ N m and nominal dc-link voltage: (**a**) evaluation-model torque and (**b**) evaluation-model voltage utilization. The electrical speed rises from 150 to 305 rad/s. The scalar comparator and verifier-gated supervisor both remain within the voltage limit but select different torque–current–margin compromises.

Table 5. Representative profile results. RMSE is in N m and MSI is in A^2; IR was 0.0% for every listed case.

Profile	Method	RMSE	VR (%)	VMI	MSI
Base transition	MTPA + rule FW	0.903	23.1	0.151	60.67
	Projection reference	0.907	0.0	0.160	61.54
	Margin-matched scalar J_ρ	1.042	0.0	0.212	61.81
	Verifier multi-candidate + CEM	1.140	0.0	0.205	59.57
Dc-link sag	MTPA + rule FW	1.152	63.7	0.001	63.61
	Projection reference	1.158	0.4	0.019	64.01
	Margin-matched scalar J_ρ	1.403	0.0	0.084	65.84
	Verifier multi-candidate + CEM	1.520	0.0	0.115	66.84
Mismatch	MTPA + rule FW	0.872	50.6	0.017	49.94
	Projection reference	0.878	50.6	0.026	50.08
	Margin-matched scalar J_ρ	0.985	0.0	0.117	56.17
	Verifier multi-candidate + CEM	0.894	0.0	0.080	53.25
Combined cycle	MTPA + rule FW	0.731	0.0	0.275	39.32
	Projection reference	0.731	0.0	0.275	39.32
	Margin-matched scalar J_ρ	0.768	0.0	0.339	42.50
	Verifier multi-candidate + CEM	0.732	0.0	0.280	39.38

Figure 3 isolates dc-link variation at $T_e^\star = 7.4$ N m and $\omega_e = 250$ rad/s. The voltage decreases during 0.30–0.36 s, remains at 278 V until 0.64 s, and recovers by 0.70 s. Figure 3a shows the resulting torque derating and recovery; Figure 3b shows that the supervisor retains more voltage headroom than the scalar comparator during this profile. The corresponding VMI values are 0.115 and 0.084, respectively. This larger reserve is accompanied by RMSE values of 1.520 versus 1.403 N m and MSI values of 66.84 versus 65.84 A^2. Both methods have zero VR, compared with 63.7% for rule-based FW and 0.4% for the projection reference. Unlike the base-speed case, the supervisor here allocates additional current and torque shortfall to voltage reserve; the common limits do not force the two decision architectures to choose the same operating point.

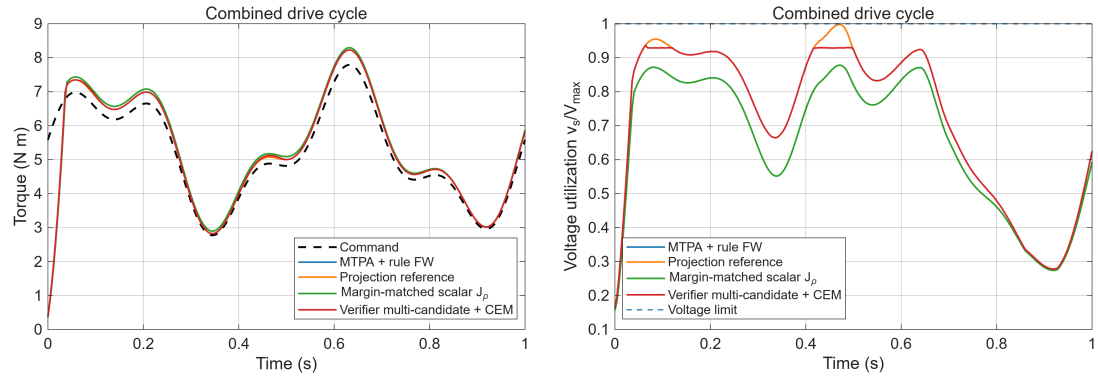


(**a**) Torque (**b**) Voltage utilization

Figure 3. Dc-link-sag profile at $T_e^\star = 7.4$ N m and $\omega_e = 250$ rad/s: (**a**) evaluation-model torque and (**b**) evaluation-model voltage utilization. The dc link falls to 278 V and then recovers. The scalar and verifier-gated methods use the same effective limits; the supervisor selects more voltage reserve in this profile.

The mismatch rows in Table 5 distinguish nominal-model feasibility from evaluation-model feasibility. Rule-based FW and nominal projection both have 50.6% evaluation-model VR, whereas the scalar comparator and supervisor maintain zero VR with the prescribed margins. Compared with the scalar comparator, the supervisor reduces RMSE from 0.985 to 0.894 N m and MSI from 56.17 to 53.25 A^2, with VMI decreasing from 0.117 to 0.080. For this specified perturbation, the multi-candidate decision therefore improves tracking and current use while retaining a smaller, but sufficient, voltage reserve.

Figure 4 examines the combined cycle with time-varying torque and speed, a temporary 292 V dc link, and the moderate parameter perturbation in Table 3. The torque curves in Figure 4a largely overlap, whereas Figure 4b separates their voltage-reserve choices. All four displayed methods have zero VR in this case. The supervisor retains near-baseline tracking and current use: its RMSE and MSI are 0.732 N m and 39.38 A^2, compared with 0.731 N m and 39.32 A^2 for the coincident rule-based and projection methods. The scalar comparator instead provides the largest average VMI, 0.339 versus 0.280 for the supervisor, with higher RMSE and MSI of 0.768 N m and 42.50 A^2. This profile demonstrates that the supervisor can retain near-baseline behavior when that behavior is admissible, rather than imposing the same degree of derating across all operating conditions.



(**a**) Torque (**b**) Voltage utilization

Figure 4. Combined drive cycle with moderate parameter mismatch and a temporary dc-link sag to 292 V: (**a**) evaluation-model torque and (**b**) evaluation-model voltage utilization. All four displayed methods remain within the voltage limit. The rule-based and projection traces coincide, while the scalar and multi-candidate methods select different current–margin compromises.

7.3. Verifier Intervention and Ranking-Weight Stress

Figure 5 reports candidate-level screening activity for the frozen CEM score. The rejection rates are 18.46% for the base-speed transition, 17.60% for the torque step, 18.59% for dc-link sag, 14.14% for parameter mismatch, and 2.39% for the combined cycle. Each percentage uses $6 \times 251 = 1506$ generated candidates as its denominator; it is not the fraction of executed references that violate a limit. The lower rejection rate in the combined cycle is consistent with the voltage-feasible traces in Figure 4, whereas the first four profiles screen a larger share of alternatives. The unweighted mean is 14.24%, as reported in Table 4. Nevertheless, score-only and gated selection produce identical option labels and current references with the calibrated weights: all rejected candidates are nonwinning alternatives. At least one candidate remains verified at every update. Thus, the gate removes inadmissible alternatives without perturbing an already admissible selected trajectory.

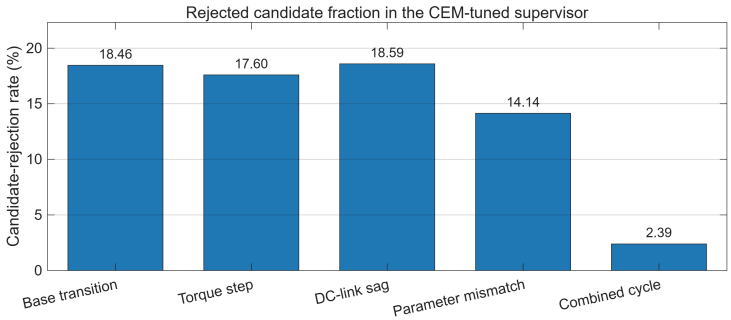


Figure 5. Candidate-rejection rate for the CEM-tuned supervisor over the five frozen evaluation profiles. Each bar uses 1506 generated candidates as its denominator. The mean rejection rate is 14.24%; the frozen-weight ablation separately confirms that the selected trajectory is unchanged by the gate.

Table 6 quantifies when the same gate becomes decisive under ranking misspecification. Each positive CEM weight is independently multiplied by a fixed-seed factor between 10^{-2} and 10^{2}, without changing the raw candidate laws, verifier conditions, or evaluation profiles. Score-only selection produces at least one controller-model current, voltage, or increment violation for 35 of 128 vectors (27.34%) and exceeds a stricter current or voltage verifier margin for 54 vectors (42.19%). The difference between these fractions distinguishes loss of the prescribed reserve from violation of a full controller-model limit. Gated selection has zero violations in both categories for every tested vector.

Table 6. Ranking-weight stress results for 128 fixed-seed vectors, with five profiles per vector and selector. The first two rows give percentages of vectors with any violation. The remaining statistics use each vector's five-profile mean metric; the VR percentile and maximum are taken across vectors. The panel is a diagnostic stress test rather than a probability distribution for deployed weights.

Metric	Score Only	Verifier Gated
Weight vectors with any controller-limit violation (%)	27.34	0.00
Weight vectors with any verifier-margin violation (%)	42.19	0.00
Mean controller-model VR (%)	5.54	0.00
95th-percentile controller-model VR (%)	34.59	0.00
Maximum controller-model VR (%)	39.60	0.00
Mean evaluation-model VR (%)	7.30	0.00
Mean torque RMSE (N m)	0.952	0.980
Mean VMI	0.178	0.191
Mean MSI (A^2)	52.43	53.01

The remaining rows of Table 6 quantify violation severity and its performance trade-off. Score-only mean controller-model VR is 5.54%, with a 95th percentile of 34.59% and a maximum of 39.60%; these last two statistics describe the 128 five-profile mean rates, not instantaneous voltage peaks. Mean evaluation-model VR is 7.30%. Both controller- and evaluation-model VR are zero for gated selection. The gate increases mean VMI from 0.178 to 0.191, with RMSE increasing from 0.952 to 0.980 N m and MSI from 52.43 to 53.01 A^2. Thus, the reduction in violations is accompanied by a quantified change in torque and current use rather than an unreported performance penalty.

Figure 6 displays three complementary summaries from Table 6. The left pair of bars measures how often a tested weight vector causes any controller-limit violation, using 128 vectors as the denominator. The center and right pairs measure average voltage-violation rates under the controller and evaluation models, respectively, across 640 rollouts per selector. Their score-only values are 27.34%, 5.54%, and 7.30%, whereas all three gated values are zero. Unlike Figure 5, which records rejected alternatives under a calibrated ranking, Figure 6 measures violations in the executed references after the ranking is perturbed. Together, the frozen-weight ablation and stress panel distinguish transparent screening from active protection and demonstrate the practical consequence of the ranking-independent admissibility property in Proposition 1.

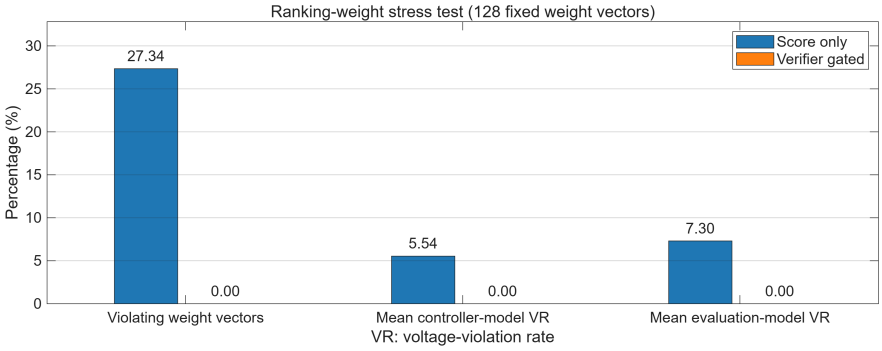


Figure 6. Ranking-weight stress results generated from the same archived records as Table 6. From left to right: the percentage of weight vectors producing any controller-limit violation, the mean controller-model voltage-violation rate, and the mean evaluation-model voltage-violation rate. The score-only values are 27.34%, 5.54%, and 7.30%; all three verifier-gated values are zero.

7.4. Candidate Contribution and Grid Sensitivity

Table 7 isolates each candidate's incremental effect with weights fixed. Across 1255 samples, MTPA, torque-preserving FW, and voltage-minimizing options are selected 322, 850, and 83 times, respectively. Projection, residual correction, and the conservative fallback are never selected in ordinary operation, and removing any one of these three leaves the reference trajectory unchanged. Removing MTPA changes only 13 samples because another candidate often supplies the same reference. No candidate is the sole verified option on this panel. The result supports a modular library, not a claim that all six options are necessary for the evaluated profiles.

Removing the torque-preserving option increases MSI from 51.488 to 53.218 A^2 and raises torque RMSE from 0.9793 to 0.9916 N m, while increasing average VMI. Its incremental role is therefore a lower-current trade-off rather than unconditional improvement of every metric. Removing the voltage-minimizing option changes 83 references and slightly lowers MSI; its contribution is also condition-dependent. In the synthetic $\pm 5\%$ box, worst utilization is 0.9591 for the full set and 0.9653 without MTPA; every ablation passes this box. The zero-effect removals remain zero-effect under fixed-reference uncertainty evaluation. These observations do not establish necessity under arbitrary drive conditions.

Table 8 tests dependence on the scalar grid without retuning the score. Changing 801 to 6401 points changes the supervisor's mean MSI by 0.00475 A^2 (0.0092%) and torque RMSE by approximately 2.0×10^{-5} N m; every tabulated grid gives zero algebraic voltage and current violations for both methods. Aggregate convergence does not imply identical reference sequences: the maximum 801-to-6401 reference difference across the five profiles is 0.2034 A for the discrete selector and 0.01213 A for the scalar comparator. This local selection sensitivity is retained in the trace-level output rather than hidden by the averages.

Table 7. Fixed-weight leave-one-candidate-out results, averaged across five profiles. Changed samples are counted over all 1255 reference updates. Emergency fallback remains available when its ordinary ranking option is removed.

Library	RMSE (N m)	MSI (A^2)	VMI	Changed Samples
Full set	0.9793	51.488	0.1822	0
Without MTPA	0.9798	51.493	0.1823	13
Without Projection	0.9793	51.488	0.1822	0
Without Torque-FW	0.9916	53.218	0.2007	444
Without Voltage-min	0.9794	51.324	0.1821	83
Without Residual	0.9793	51.488	0.1822	0
Without Fallback	0.9793	51.488	0.1822	0

Table 8. Grid-resolution comparison with frozen controller settings. Columns report five-profile means; RMSE is in N m and MSI in A^2. The 6401-point result is a refinement reference, not an exact solution.

Points	CEM RMSE	CEM MSI	Scalar RMSE	Scalar MSI
201	0.979482	51.35080	0.937522	53.34194
401	0.979412	51.45990	0.937728	53.34220
801	0.979339	51.48750	0.937646	53.33168
1601	0.979327	51.48523	0.937638	53.33918
3201	0.979325	51.48298	0.937629	53.33662
6401	0.979319	51.48275	0.937629	53.33557

7.5. Parameter Envelopes and Torque Capability

Table 9 locates where the fixed verifier reserve is insufficient under parameter uncertainty. Proposition 3 makes the vertex maximum exact over each specified box for a fixed reference, apart from numerical arithmetic. Both reference sequences meet the full voltage limit throughout the $\pm 5\%$ box. At $\pm 10\%$, the supervisor's worst utilization is 1.0159, and the scalar comparator's is 1.0547; neither method then has a voltage guarantee over that box. At $\pm 20\%$, the peaks rise to 1.1341 and 1.1767. Thus, ranking independence does not make the nominal margin independent of motor-parameter error.

Bisection of a common symmetric box gives limiting voltage-envelope radii of approximately 8.64% for the supervisor and 5.46% for the scalar reference panel. These values apply only to the 1255 fixed references and operating samples, not all operating points or uncertainty-dependent closed-loop trajectories. The violation-duration column in Table 9 is the largest descriptive profile rate among tested vertices; it is not a probability of failure. The capability-capped metric further separates the chosen reserve from unavailable command torque. After the first 100 ms, the five-profile mean algebraic shortfall relative to the grid-estimated achievable command is 0.2399 N m for the supervisor, 0.1306 N m for the scalar comparator, and 0.0008 N m for nominal projection. Projection's small shortfall must be read together with its nonzero voltage violations in the primary comparison. The supervisor's lower MSI therefore does not mean greater torque capability; part of the current reduction is accompanied by additional torque derating.

Table 9. Fixed-reference parameter-box results. Worst utilization is the maximum over all vertices and samples. Maximum profile VR is the largest profile violation percentage among vertices, not a reliability estimate.

Box Radius	Method	Worst Voltage (p.u.)	Maximum Profile VR (%)
±5%	Supervisor	0.9591	0.00
±5%	Scalar	0.9945	0.00
±10%	Supervisor	1.0159	95.22
±10%	Scalar	1.0547	96.02
±20%	Supervisor	1.1341	96.81
±20%	Scalar	1.1767	96.41

7.6. Electrical Closed-Loop Transients

Table 10 reports every nonideal switching case. These are integrated plant-current and torque results, rather than algebraic outputs evaluated at the reference. No 10 A current-limit violation occurs in any of the 54 primary runs or 6 additional continuous-optimizer runs. In the six nonideal switching cases, the supervisor's largest current magnitude is 8.819 A. This finite-panel result does not replace a current-tracking error bound or a physical-drive safety certificate.

Table 10. Nonideal switching results with delayed current regulation, 2 µs dead time, 0.8 V device drop, and 0.02 A sensing-noise standard deviation. RMSE is in N m and MSI in A^2. RMSE includes startup; current peaks are sampled at 1 µs. $u_{\mathrm{dem,pk}}$ is peak pre-saturation demand; Sat. is the percentage of saturated current-loop periods.

Profile	Method	RMSE	MSI	Peak A	$u_{\mathrm{dem,pk}}$	Sat. (%)
Base ramp	Projection	1.046	56.00	7.846	1.134	37.80
Base ramp	Scalar	1.061	61.68	8.516	1.106	0.59
Base ramp	Supervisor	1.157	59.49	8.558	1.088	0.22
Base ramp	Direct NLP	1.153	56.79	7.784	1.002	0.01
Smooth torque step	Projection	0.487	32.90	7.668	1.326	31.36
Smooth torque step	Scalar	0.503	40.20	7.871	1.311	0.35
Smooth torque step	Supervisor	0.636	38.00	7.631	1.270	0.45
Smooth torque step	Direct NLP	0.650	36.32	7.509	1.069	0.53
DC-link sag	Projection	1.459	53.51	8.150	1.390	80.08
DC-link sag	Scalar	1.420	65.50	8.906	1.192	1.68
DC-link sag	Supervisor	1.545	66.19	8.819	1.281	0.61
DC-link sag	Direct NLP	1.540	60.06	8.389	1.147	0.47
Parameter mismatch	Projection	1.164	42.78	7.507	1.088	57.30
Parameter mismatch	Scalar	0.994	55.44	7.816	1.084	8.56
Parameter mismatch	Supervisor	0.908	52.88	7.719	1.088	4.39
Parameter mismatch	Direct NLP	0.829	49.52	7.453	1.093	5.72
Combined cycle	Projection	0.770	38.95	8.377	1.080	3.57
Combined cycle	Scalar	0.808	42.40	8.525	1.026	0.03
Combined cycle	Supervisor	0.772	39.21	8.384	1.080	0.51
Combined cycle	Direct NLP	0.780	38.68	8.284	1.026	0.25
Hard torque step	Projection	0.547	33.99	7.668	1.256	34.34
Hard torque step	Scalar	0.574	40.79	7.877	1.113	0.30
Hard torque step	Supervisor	0.725	38.21	7.632	1.281	0.64
Hard torque step	Direct NLP	0.726	36.96	7.505	1.164	0.58

The trade-off remains visible dynamically. During dc-link sag, the supervisor has 1.545 N m torque RMSE versus 1.420 N m for the scalar comparator, and MSI is also slightly higher. During parameter mismatch, the supervisor instead reduces torque RMSE to 0.908 N m versus 0.994 N m and reduces MSI. Its saturation duration is 4.39% versus 8.56% in that case. The results identify useful operating conditions without assigning a universal winner. The direct NLP gives 0.829 N m RMSE in the mismatch case and 60.06 A^2 MSI

during sag, compared with the supervisor's 0.908 N m and 66.19 A^2, respectively. Its hard-step RMSE is 0.726 N m, essentially the same as the supervisor's 0.725 N m; continuous optimization is not a uniform dynamic tracking improvement either.

Figure 7 shows how the integrated torque responds to an imposed speed ramp and a discontinuous torque command. In the ramp, voltage-limited torque separates the four reference methods after the common startup transient. In the hard step, the scalar reference tracks the higher command more closely than the supervisor, consistent with Table 10; neither the plot nor the RMSE conceals deliberate derating as improved tracking. These curves complement, rather than relabel, the quasi-steady trajectories in Figures 2–4. Figure 8 separates controller voltage demand from actual current tracking during the sag. Projection requests voltage above the saturation threshold for a substantial part of the interval, whereas the margin-based references shorten this saturation exposure. The supervisor's d- and q-axis states follow the issued references with finite error and switching ripple; they are not identified algebraically with those references. This distinction is the reason the model-side theorem and the current-loop evidence are reported separately.

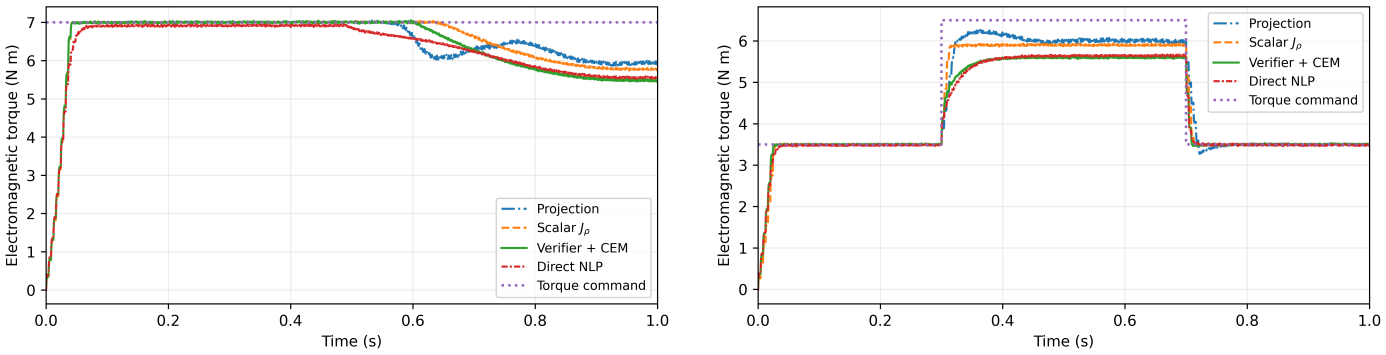


Figure 7. Integrated electromagnetic torque under the nonideal switching model: (**a**) imposed electrical-speed ramp; (**b**) hard torque step at fixed electrical speed. The dotted curve is the requested torque; projection, scalar, supervisor, and direct-NLP responses use the same current controller and inverter settings.

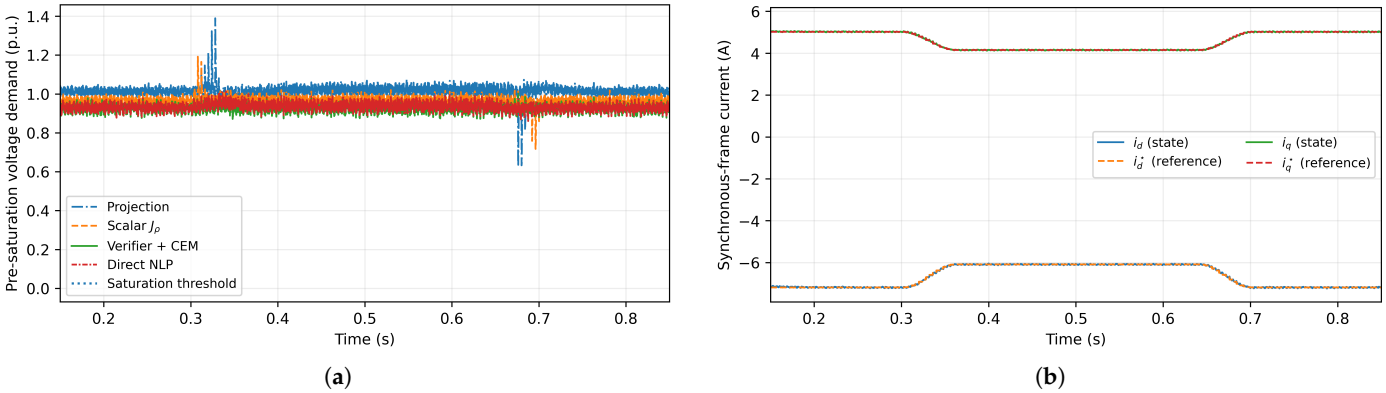


Figure 8. Nonideal switching dc-link-sag detail: (**a**) pre-saturation fundamental-voltage demand for the four methods; (**b**) supervisor current states and references. The horizontal threshold in (**a**) identifies commanded-voltage saturation, not an instantaneous switching-vector limit.

Halving the integration step to 0.5 µs in the two prespecified sag cases changes torque RMSE by less than 9.1×10^{-5} N m. RMS differences in sampled current trajectories are below 0.0089 A, although the largest pointwise difference is 0.0449 A. The independent matrix-exponential check of the electrical equations gives errors below 1.7×10^{-13} A in

the three constant-input test cases. These checks address numerical implementation and resolution, not uncertainty in the assumed physical model.

7.7. Emergency States and Host Cost

Table 11 exercises states absent from the smooth primary profiles. Abrupt reductions to 150, 100, and 50 V reject all rate-limited candidates but retain an electrically feasible full fallback. The returned increments exceed 0.75 A and are explicitly flagged as status 1; Proposition 2 therefore does not apply at those updates. At 5 V, the computed fallback does not pass the electrical check and the interface returns status 2 rather than executing an unchecked zero-current point. The diagnostic exposes both rate-priority conflict and lack of a certified output.

Table 11. Single-update emergency tests from a common preceding operating reference at 7.4 N m and 250 rad/s. Status 1 explicitly overrides the reference-increment cap. Status 2 returns no actuator reference.

DC Link (V)	Status	Returned Outcome	Increment (A)
150	1	Electrical fallback/override	4.259
100	1	Electrical fallback/override	5.311
50	1	Electrical fallback/override	6.270
5	2	No certified reference	—

The emergency path requires at most one additional electrical fallback check after the six ordinary candidate checks. The convex-segment argument in Section 4.3 explains why a backward scan of 101 points is redundant when the raw fallback is electrically feasible. Checking the fallback itself is essential because feasibility cannot be inferred merely from a routine returning a numeric pair. The single-update interface reproduces all five primary supervisor trajectories exactly, so the exceptional-state handling does not change their reported metrics.

For a grid of G points, shared candidate construction uses two main grid passes and at most five fallback passes, followed by constant-size six-candidate processing. Its cost is $O(G)$, with at most $7G$ grid-loop iterations in this implementation; individual iterations have different arithmetic work and are not CPU cycles. The grid, six raw current targets, six score weights, and previous reference/label account for 6576 bytes at $G = 801$. This subtotal excludes intermediate arrays, stack, compiled code, and runtime. Peak resident memory for the full Python/JIT audit process is 236,404 KiB, not the memory requirement of a deployed controller.

Table 12 reports host wall-clock observations, retaining the large scalar outlier. The scalar path has a lower median latency than the supervisor, whereas maxima are sensitive to occasional long observations. These samples characterize this host implementation only. They do not establish a target-DSP WCET or guarantee completion before every 4 ms deadline.

Table 12. Single-update host wall-clock measurements. Compilation and warm-up are excluded; no observed outlier is removed. The maximum column is in milliseconds; other latency columns are in microseconds. These are neither target-controller timing nor certified WCET.

Path	Calls	Median (µs)	95th (µs)	99th (µs)	Max. (ms)
Supervisor	12,550	61.92	283.10	563.72	2.890
Supervisor emergency	4000	45.01	185.31	389.68	2.511
Scalar	12,550	14.08	78.88	228.71	486.234
Direct NLP, input panel	12,550	1457.670	3110.882	5223.010	144.487

7.8. Matched Calibration and Broader Optimization Comparison

Table 13 evaluates the selected settings using the frozen five-profile panel. Equal-budget calibration reduces the supervisor's development cost from 0.033476 to 0.031044, the specified feedback variant's cost from 0.556081 to 0.104587, and the scalar schedule's cost from 0.038237 to 0.038163. The selected supervisor ratios, in the order of (23), are $[1, 0.057512, 0.165595, 0.076351, 0.016440, 1.328679]$. The selected feedback gains are $K_P = 40.8899\ \mathrm{A/pu}$ and $K_I = 11.2682\ \mathrm{A/(pu\,s)}$; the scalar transition speed is $\omega_a = 228.614\ \mathrm{rad/s}$. Full-precision values and all 483 cost evaluations are archived. The primary ablation, uncertainty, dynamic-supervisor, and supervisor-timing results retain the frozen calibration. Matched-tuning variants are reported separately in Table 13.

The matched supervisor retains lower MSI than the matched scalar (50.899 versus 52.979 A^2) at higher torque RMSE (0.9793 versus 0.9344 N m). The explicitly defined feedback variant still has 23.90% mean evaluation-model voltage violations; this is a finding about that implementation and search domain, not every published voltage-feedback controller. Equal tuning budgets remove one comparison imbalance without implying that unrelated control architectures are identical.

The direct NLP obtains both lower MSI (48.440 A^2) and lower torque RMSE (0.9624 N m) than the frozen CEM supervisor. The supervisor nevertheless has slightly lower common cost, 0.04941 versus 0.05074. Component-wise recomputation identifies the voltage-penalty term as the cause: it averages 0.00871 for the supervisor and 0.01200 for the NLP, whereas the NLP has smaller tracking, current, and increment terms. Consequently, neither the primary scalar comparison nor the shared objective establishes a universal performance advantage for multi-candidate selection. Every one of the 2259 development/evaluation NLP updates returns an independently verified reference. Three of 4518 individual solver attempts have non-success status, which is retained in the logs; a solver flag alone is not used as a feasibility certificate. The complete online NLP path, including its fallback initializer, has 1457.67 µs median host latency, compared with 61.92 µs for the supervisor on the same operating-input/history panel. Its observed maximum is 144.49 ms. The host/software-specific comparison quantifies one computational trade-off of direct continuous search, without converting a sampled maximum into WCET or guaranteeing a 4 ms deadline.

Table 13. Five-profile means under the common criterion $\mathcal{J}_c$. Matched tuning uses 161 development-cost calls per tunable method. All rows have zero current violations; voltage violations remain visible. The direct NLP uses a fixed shared one-step criterion, not fitted score weights.

Method	$\mathcal{J}_c$	RMSE (N m)	VR (%)	MSI (A^2)	VMI
Rule, fixed	1.20472	0.7937	34.74	49.700	0.1175
Projection, fixed	0.43298	0.7969	10.20	50.603	0.1278
Feedback, matched tuning	0.68176	0.8176	23.90	54.017	0.1321
Scalar, fixed	0.05630	0.9376	0.00	53.332	0.2010
Scalar, matched tuning	0.05653	0.9344	0.00	52.979	0.1987
Supervisor, frozen CEM	0.04941	0.9793	0.00	51.488	0.1822
Supervisor, matched tuning	0.04932	0.9793	0.00	50.899	0.1753
Direct NLP, fixed criterion	0.05074	0.9624	0.00	48.440	0.1763

8. Discussion

The results establish two complementary properties of the proposed architecture. First, Proposition 1 makes model-side admissibility independent of the ranking weights whenever the verified set is nonempty. Second, the numerical study quantifies the performance cost and benefit of exercising that architecture. Across the five frozen profiles, both the margin-matched scalar comparator and the verifier-gated supervisors maintained zero

evaluation-model voltage and current violations, while the CEM supervisor used 3.46% less mean squared current than the scalar comparator. The corresponding cost was 4.4% higher torque RMSE and 0.019 lower average voltage headroom. These differences are consistent with their decision structures: the scalar method continuously minimizes one scheduled objective, whereas the supervisor chooses among six physically distinct laws using torque, voltage, current, reference-change, option-switching, and margin terms.

The ranking-weight analysis isolates the verifier's architectural effect. With calibrated CEM weights, the gate screened nonwinning candidates and left the selected trajectory unchanged, so nominal behavior was determined by the candidate set and score. Under broad weight misspecification, the same gate changed the outcome precisely when required: 35 of 128 score-only weight vectors authorized at least one controller-limit violation, whereas all gated weight vectors remained admissible under the encoded model. The mean controller-model VR fell from 5.54% to 0, and the mean evaluation-model VR fell from 7.30% to 0 across 640 rollouts. Together, Proposition 1 and the stress results establish ranking-independent model-side admissibility while allowing the ranking to express different performance priorities.

The closest safeguarded-scalar framework [26] supplies objective-independent continuous feasibility and optimization. The contribution here is the discrete orchestration layer that allows heterogeneous analytical, projection, residual-corrected, or product-specific reference laws to coexist without analytically merging them into one convex objective. Previous-reference and previous-option terms can be introduced directly, and the ranking can be retuned or replaced without altering the admissibility conditions. The six implemented options provide a modular library, while the fixed-weight ablation shows that only three labels are selected on this particular panel. Candidate diversity is an architectural capability, not evidence that every included option improves every operating condition.

CEM provides a reproducible offline calibration mechanism rather than the safety property itself. It reduced the development objective by 7.28% and increased held-out VMI by 0.007 with essentially unchanged torque RMSE while modestly increasing MSI and MSR. The same verifier property would hold with manual weights, another optimizer, or a rule-based ranking policy.

The parameter-envelope analysis and electrical closed-loop tests address different gaps between nominal reference feasibility and drive behavior. The former certifies the voltage calculation over a specified box for fixed references; the latter follows actual simulated current states under a delayed, saturating current loop and selected inverter effects. Their results agree that the verifier's ranking independence is not a universal robustness guarantee. Parameter boxes outside the available reserve expose voltage failures, and ordinary margin-based references can still produce transient controller saturation.

The evidence remains specific to one constant-inductance benchmark and synthetic nonideality settings. Nonlinear magnetic saturation maps, temperature-dependent flux and inductance surfaces, mechanical speed-loop interaction, real-time hardware-in-the-loop, and physical drive tests are outside this validation. A status-2 output also requires a separately designed protection interface. The host timing and memory inventory establish implementation costs at a stated software level, not embedded deadline or physical protection guarantees.

9. Conclusions

This paper developed a verifier-gated multi-candidate supervisor whose encoded current, voltage-margin, branch, denominator, and reference-increment constraints are independent of the ranking weights. Six physics-based current-reference laws are coordinated by a history-dependent score, while the deterministic gate retains authority over admissi-

bility. Across five frozen quasi-steady profiles, the three conventional non-verifier baselines produced 10.2–38.3% average voltage-violation rates, whereas the margin-matched scalar comparator and both gated supervisors produced 0.0% voltage and current violations. Relative to the scalar comparator, the CEM supervisor reduced MSI from 53.33 to 51.49 A^2 at the cost of increasing RMSE from 0.938 to 0.979 N m and reducing VMI from 0.201 to 0.182. The gate was transparent under the calibrated weights but decisive under weight misspecification: it reduced violating weight vectors from 35 of 128 to 0, mean controller-model VR from 5.54% to 0, and evaluation-model VR from 7.30% to 0 across 640 stress rollouts. The diagnostic and dynamic results delimit where this architectural property is useful. Fine-grid changes have little effect on aggregate metrics, while candidate ablation exposes redundant options in the chosen profiles. The synthetic $\pm 5\%$ parameter box preserves fixed-reference voltage feasibility; larger boxes do not. All 60 electrical closed-loop simulations remain below the 10 A current limit, with the existing torque–current trade-off retained. Exceptional outputs are explicitly separated into electrical fallback with a rate override and no certified reference. The continuous optimizer uses less current and has lower aggregate torque error, while the finite selector has lower median host latency in this implementation. The contribution is therefore model-side admissibility with transparent numerical costs and limits, rather than an unconditional physical-drive or real-time guarantee.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/electronics15214954/s1: MATLAB scripts for Figures 2–6.

Author Contributions: Conceptualization, H.H.C.; methodology, H.H.C.; software, H.H.C.; validation, T.K. and H.H.C.; formal analysis, H.H.C.; investigation, T.K. and H.H.C.; writing—original draft preparation, H.H.C.; writing—review and editing, T.K. and H.H.C. All authors have read and agreed to the published version of the manuscript.

Funding: The authors are thankful to human resources development project of the Korea Institute of Energy Technology Evaluation and Planning (KETEP) for the support grant funded by the Korea government Ministry of Trade, Industry, and Energy under the project titled "Middle market enterprise specialized human resources development for residential and commercial fuel cells", numbered: 20224000000580.

Data Availability Statement: The numerical data reported in the tables are contained within the article. The Base-MATLAB script, compatible with MATLAB R2024a or later, used to reproduce Figures 2–6 and generate their associated plotting data is provided as Supplementary Material.

Acknowledgments: During preparation of this manuscript, the authors used ChatGPT 5.6 (OpenAI) for language editing, manuscript organization, and assistance in drafting reproducibility scripts. The authors independently verified the equations, code, numerical outputs, references, and final text and take full responsibility for the content.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CEM	cross-entropy method
FOC	field-oriented control
FW	flux weakening
IPMSM	interior permanent magnet synchronous motor
IR	current-violation rate
MSI	mean-squared-current index
MSR	mode-switching ratio
MTPA	maximum torque per ampere

MTPV	maximum torque per voltage
PMSM	permanent magnet synchronous motor
RMSE	root mean square error
RR	candidate-rejection rate
VMI	voltage-margin index
VR	voltage-violation rate

References

1. Krause, P.C.; Wasynczuk, O.; Sudhoff, S.D.; Pekarek, S. *Analysis of Electric Machinery and Drive Systems*, 3rd ed.; Wiley-IEEE Press: Hoboken, NJ, USA, 2013.
2. Sul, S.-K. *Control of Electric Machine Drive Systems*; Wiley-IEEE Press: Hoboken, NJ, USA, 2011.
3. Bose, B.K. *Modern Power Electronics and AC Drives*; Prentice Hall: Upper Saddle River, NJ, USA, 2002.
4. Jahns, T.M.; Kliman, G.B.; Neumann, T.W. Interior permanent-magnet synchronous motors for adjustable-speed drives. *IEEE Trans. Ind. Appl.* **1986**, *IA-22*, 738–747. [CrossRef]
5. Jahns, T.M. Flux-weakening regime operation of an interior permanent-magnet synchronous motor drive. *IEEE Trans. Ind. Appl.* **1987**, *IA-23*, 681–689. [CrossRef]
6. Macminn, S.R.; Jahns, T.M. Control techniques for improved high-speed performance of interior PM synchronous motor drives. *IEEE Trans. Ind. Appl.* **1991**, *27*, 997–1004. [CrossRef]
7. Morimoto, S.; Takeda, Y.; Hirasa, T.; Taniguchi, K. Expansion of operating limits for permanent magnet motor by current vector control considering inverter capacity. *IEEE Trans. Ind. Appl.* **1990**, *26*, 866–871. [CrossRef]
8. Morimoto, S.; Sanada, M.; Takeda, Y. Wide-speed operation of interior permanent magnet synchronous motors with high-performance current regulator. *IEEE Trans. Ind. Appl.* **1994**, *30*, 920–926. [CrossRef]
9. Kim, J.-M.; Sul, S.-K. Speed control of interior permanent magnet synchronous motor drive for the flux-weakening operation. *IEEE Trans. Ind. Appl.* **1997**, *33*, 43–48. [CrossRef]
10. Lu, D.; Kar, N.C. A review of flux-weakening control in permanent magnet synchronous machines. In *Proceedings of the 2010 IEEE Vehicle Power and Propulsion Conference, Lille, France, 1–3 September 2010*; IEEE: New York, NY, USA, 2010; pp. 1–6.
11. Jing, R.; Wang, G.; Zhang, G.; Xu, D. Review of field weakening control strategies of permanent magnet synchronous motors. *CES Trans. Electr. Mach. Syst.* **2024**, *8*, 319–331. [CrossRef]
12. Alwaqfi, S.; Alzayed, M.; Chaoui, H. Flux-weakening control methods for permanent magnet synchronous machines in electric vehicles at high speed. *Electronics* **2025**, *14*, 3779. [CrossRef]
13. Uddin, M.N.; Radwan, T.S.; Rahman, M.A. Performance of interior permanent magnet motor drive over wide speed range. *IEEE Trans. Energy Convers.* **2002**, *17*, 79–84. [CrossRef]
14. Kwon, T.; Sul, S.-K. Novel antiwindup of a current regulator of a surface-mounted permanent-magnet motor for flux-weakening control. *IEEE Trans. Ind. Appl.* **2006**, *42*, 1293–1300. [CrossRef]
15. Sue, S.-M.; Pan, C.-T. Voltage-constraint-tracking-based field-weakening control of IPM synchronous motor drives. *IEEE Trans. Ind. Electron.* **2008**, *55*, 340–347. [CrossRef]
16. Rahman, M.F.; Zhong, L.; Lim, K.W. A direct torque-controlled interior permanent magnet synchronous motor drive incorporating field weakening. *IEEE Trans. Ind. Appl.* **1998**, *34*, 1246–1253. [CrossRef]
17. Stumper, J.-F.; Dötlinger, A.; Jung, J.; Kennel, R. Predictive control of a permanent magnet synchronous machine based on real-time dynamic optimization. In *Proceedings of the 2011 14th European Conference on Power Electronics and Applications, Birmingham, UK, 30 August–1 September 2011*; IEEE: New York, NY, USA, 2011; pp. 1–8.
18. Stumper, J.-F.; Dötlinger, A.; Kennel, R. Classical model predictive control of a permanent magnet synchronous motor. *EPE J.* **2012**, *22*, 24–31. [CrossRef]
19. Pellegrino, G.; Armando, E.; Guglielmi, P. Direct-flux vector control of IPM motor drives in the maximum torque per voltage speed range. *IEEE Trans. Ind. Electron.* **2012**, *59*, 3780–3788. [CrossRef]
20. Eldeeb, H.; Hackl, C.M.; Horlbeck, L.; Kullick, J. A unified theory for optimal feedforward torque control of anisotropic synchronous machines. *Int. J. Control* **2018**, *91*, 2273–2302. [CrossRef]
21. Bianchini, C.; Bisceglie, G.; Torreggiani, A.; Davoli, M.; Macrelli, E.; Bellini, A.; Frigieri, M. Effects of the magnetic model of interior permanent magnet machine on MTPA, flux weakening and MTPV evaluation. *Machines* **2023**, *11*, 77. [CrossRef]
22. Miguel-Espinar, C.; Heredero-Peris, D.; Gross, G.; Llonch-Masachs, M.; Montesinos-Miracle, D. Maximum torque per voltage flux-weakening strategy with speed limiter for PMSM drives. *IEEE Trans. Ind. Electron.* **2021**, *68*, 9254–9264. [CrossRef]
23. Bae, M.; Kim, S.-M.; Choi, H.H. Explicit algebraic approximations for MTPA, MTPV, and loss-minimization optimal control of PMSMs. *Electronics* **2026**, *15*, 1440. [CrossRef]

24. Kim, S.-M.; Choi, H.H. Low-complexity residual-corrected loss-minimization current-reference generation for PMSM drives. *Electronics* **2026**, *15*, 3000. [CrossRef]
25. Kang, D.; Choi, H.H. Low-complexity SVS current-reference generation for PMSMs with and without MTPV endpoints. *Energies* **2026**, *19*, 3621. [CrossRef]
26. Kang, D.; Choi, H.H. A unified safeguarded scalar framework for convex performance-index-based current-reference generation in IPMSM drives under current and voltage constraints. *Electronics* **2026**, *15*, 4313. [CrossRef]
27. Rubinstein, R.Y. The cross-entropy method for combinatorial and continuous optimization. *Methodol. Comput. Appl. Probab.* **1999**, *1*, 127–190. [CrossRef]
28. de Boer, P.T.; Kroese, D.P.; Mannor, S.; Rubinstein, R.Y. A tutorial on the cross-entropy method. *Ann. Oper. Res.* **2005**, *134*, 19–67. [CrossRef]
29. SciPy Community. minimize(method='SLSQP'). SciPy Reference Guide. Available online: https://docs.scipy.org/doc/scipy/reference/optimize.minimize-slsqp.html (accessed on 17 October 2026).
30. Fan, Y.; Ma, H.; Zhu, G.; Luo, J. Improved MTPA and MTPV optimal criteria analysis based on IPMSM nonlinear flux-linkage model. *Energies* **2024**, *17*, 3494. [CrossRef]

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