

Article

A Unified Safeguarded Scalar Framework for Convex Performance-Index-Based Current-Reference Generation in IPMSM Drives Under Current and Voltage Constraints

Dongyeop 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; 2020111876@dgu.ac.kr

* Correspondence: hhchoi@dongguk.edu; Tel.: +82-2-2260-3777

Abstract

Interior permanent-magnet synchronous motor (IPMSM) current-reference laws are commonly derived for one objective or operating region at a time. This study develops a safeguarded scalar framework that separates exact torque-equality elimination, current/voltage feasibility, torque-command derating, and performance-index selection. On the positive-motoring branch, the torque equality reduces the two-current problem to one variable, while convex current and voltage residuals define a connected feasible interval. Taylor-based quadratic boundary estimates are refined by fixed-count Newton steps and accepted only after residual, slope, and interval checks; failed candidates are routed to a deterministic convex-sublevel fallback. For a strictly convex reduced objective, safeguarded Newton proposals are followed by mandatory derivative-bisection contractions, which provide an explicit final-bracket error bound. Maximum-torque-per-ampere, fixed-torque voltage minimization, loss-oriented indices, a normalized current–voltage stress index, and two structurally convex design extensions are treated in the same pipeline. A 3844-point steady-state numerical sweep over two published IPMSM parameter sets and both zero-resistance and full-resistance voltage models contained 3071 feasible cases. The low-cost setting produced a maximum d -axis current error of 1.331×10^{-3} A, whereas the conservative setting reduced it to 2.892×10^{-4} A. The evidence is numerical and steady-state; processor timing and hardware validation remain future work.

Keywords: interior permanent-magnet synchronous motor; current-reference generation; feasible interval; convex optimization; MTPA; MTPV; loss-minimization control; field weakening; safeguarded Newton method; derivative bisection



Academic Editors: Chengbo Yang,
Yuanlong Xie and Mingjin Hu

Received: 18 August 2026

Revised: 14 September 2026

Accepted: 18 September 2026

Published: 20 September 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](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Permanent-magnet synchronous motor (PMSM) drives use a current-reference generator to convert a torque command into a pair $(\hat{i}_d, \hat{i}_q)$ that is compatible with current and voltage limits. For an interior PMSM (IPMSM), saliency provides an additional degree of freedom: several current pairs can generate the same torque, and the selected pair can be used to reduce current magnitude, voltage utilization, modeled loss, or another engineering index. Classical maximum-torque-per-ampere (MTPA), field-weakening (FW), maximum-torque-per-voltage (MTPV), and loss-minimization control (LMC) laws address these goals from different perspectives [1–5].

The relevant literature includes analytical, numerical, data-driven, and loss-oriented approaches. Ferrari-based solutions expose the quartic structure of current-minimizing

IPMSM torque control [6], whereas unified analytical theory covers optimal feedforward torque control of anisotropic synchronous machines with non-negligible stator resistance and mutual inductance [7]. Experimentally validated numerical schemes span MTPA, FW, maximum-current, and MTPV regions for the minimum-current objective [8]. Recent reviews classify calculation-based, voltage-loop, and predictive FW methods and emphasize their different trade-offs in model dependence, robustness, calibration, and implementation effort [9,10]. Recent MTPA methods are surveyed in [11]; parameter-aware tracking has been addressed through sparse-Bayesian offline mapping with online magnet-temperature compensation [12], and numerical MTPA optimization explicitly considers nonlinear flux linkage, saturation, and cross-coupling [13]. Efficiency-oriented studies have likewise incorporated iron-loss resistance, bivariate-extreme-value optimization, linearized torque/loss models, and drive-system-level equivalent-circuit optimization [14–17]. These developments show that the selected performance objective and the constraint-handling architecture are related but distinct design layers. Advanced PMSM control also includes composite adaptive super-twisting sliding-mode control using a barrier function [18] and a super-twisting extended-state-observer-based quasi-proportional-resonant controller [19]. These methods address dynamic regulation and disturbance rejection in the surrounding drive. The present work supplies the constrained current-reference optimization layer; compatibility with a particular feedback controller still requires closed-loop analysis and validation.

Two recent publications from the authors' group are the closest antecedents. Kim and Choi [20] reduced one iron-loss-based LMC objective to a scalar stationarity residual and applied a one-step Newton-type correction to an explicit warm start. Kang and Choi [21] developed a speed- and voltage-margin scheduled anchor blend, followed by torque-consistent projection and derating, for cost-sensitive wide-speed-range drives. In contrast, the present study develops a common feasible-interval and scalar-optimization interface for selectable convex performance indices, with certified boundary candidates, monotone positive-torque feasibility, and a derivative-sign bracket that yields an explicit bound on the scalar objective-search error.

The remaining implementation question is whether torque equality and constraint feasibility can be handled once, after which certified convex objectives can be inserted without re-deriving region-specific constraint equations. A dependable scalar formulation must also reject inaccurate fixed-count Newton candidates and preserve a certified objective bracket. The proposed method therefore proves the interval structure, establishes monotone torque feasibility for the positive-motoring branch, removes the zero-speed singularity from the simplified voltage residual, certifies Taylor/Newton boundary candidates, and applies deterministic derivative-bracket contraction.

Torque elimination, convex-sublevel geometry, and bisection are established mathematical ingredients. The algorithmic contribution here is their explicit assembly into four reusable modules: (i) current-first interval construction with a voltage-inactive shortcut; (ii) branch-oriented Taylor/Newton acceptance, a feasible-anchor check, and bounded boundary fallback; (iii) an objective-registration interface that separates proven convex indices from general scalar objectives; and (iv) optional Newton acceleration followed by mandatory bracket contraction, with feasibility and derating status exposed to the supervisor. The positive-torque nesting proof and nonsingular zero-speed voltage residual establish the conditions under which these modules can be composed. Throughout this work, a numerically “certified” candidate means that it has passed the stated finite-precision residual, derivative-orientation, domain, and feasibility checks, with deterministic fallback on rejection. This is an operational numerical qualification, not interval-arithmetic certification or a proof of physical drive feasibility under uncertain parameters. Ana-

lytical convexity and bracket theorems are stated separately under their mathematical assumptions. Section 8.12 examines the numerical classification limits.

The central claim is the following: under the stated constant-parameter, steady-state IPMSM assumptions, constrained current-reference generation for a positive motoring torque can be organized as torque elimination $\rightarrow$ feasible i_d interval $\rightarrow$ safeguarded scalar objective minimization with objective-independent feasibility handling and an explicit fallback whenever a fast candidate is not certified.

The specific contributions are as follows: first, an objective-independent feasible-interval formulation that provides a common constraint interface for selectable performance indices; second, a current-first interval construction, a zero-torque feasibility precheck, and a proof that positive-torque feasibility is monotone, which justifies scalar torque bisection; third, a certified Taylor/Newton boundary path combined with a fixed-count convex-sublevel fallback that avoids online quartic formulas without accepting unchecked roots; fourth, a safeguarded objective solver in which optional Newton proposals preserve a derivative-sign bracket and mandatory derivative-bisection contractions provide an a priori bound on the scalar objective-search error within the certified feasible interval; fifth, reproducible numerical verification over representative points, a 31×31 speed–torque grid for each motor and voltage model, and a fixed-seed randomized stress test. This paper additionally evaluates MTPA and voltage minimization, magnetic-parameter mismatch, dc-link variation, tolerance sensitivity, and the full residual/derivative evaluation burden. The analysis covers a steady-state, constant-parameter IPMSM with $L_q > L_d$, nonnegative electrical speed, positive motoring torque, and the practical branch $i_d \leq 0$, $i_q \geq 0$. Section 8.14 discusses the extensions required for saturation, cross-coupling, inverter nonlinearities, current-control dynamics, reference-transition smoothing, negative torque, and processor-specific timing.

2. IPMSM Model and Constrained Current-Reference Problem

2.1. Steady-State Model

In the synchronously rotating dq reference frame, the steady-state voltage equations are

$$v_d = R_s i_d - \omega L_q i_q, \quad (1)$$

$$v_q = R_s i_q + \omega (\lambda_m + L_d i_d), \quad (2)$$

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 is the electrical angular speed. The electromagnetic torque is

$$T_e = K_t \{ \lambda_m + (L_d - L_q) i_d \} i_q, \quad (3)$$

where $K_t = 3p/4$ when p denotes the number of poles; equivalently, $K_t = 3p_p/2$ for p_p pole pairs.

The current and voltage constraints are

$$i_d^2 + i_q^2 \leq I_M^2, \quad (4)$$

$$v_d^2 + v_q^2 \leq V_M^2. \quad (5)$$

Here, I_M and V_M denote the maximum allowable stator-current and stator-voltage magnitudes, respectively. For a selected performance index $J(i_d, i_q)$ and a torque command T_r , the current-reference problem is

$$\begin{aligned} \min_{i_d, i_q} \quad & J(i_d, i_q) \\ \text{subject to} \quad & T_e = T_r, \quad i_d^2 + i_q^2 \leq I_M^2, \quad v_d^2 + v_q^2 \leq V_M^2, \quad i_d \in \mathcal{I}_{\text{op}}. \end{aligned} \quad (6)$$

Assumption 1 (Operating branch and model scope). *The analysis uses $T_r \geq 0$, $\omega \geq 0$, $L_q > L_d > 0$, $\lambda_m > 0$, and $R_s \geq 0$. The practical search interval is*

$$\mathcal{I}_{\text{op}} = [i_{d,L}^{\text{op}}, 0], \quad i_{d,L}^{\text{op}} = \max(-I_M, i_{d,\text{demag}}), \quad (7)$$

where $i_{d,\text{demag}}$ is an optional demagnetization bound. The selected branch satisfies $i_d \leq 0$ and $i_q \geq 0$. Motor parameters are treated as constant at each evaluated operating point.

2.2. Exact Torque Elimination

For a given T_r , define

$$C = \frac{T_r}{K_t}, \quad \Delta = L_d - L_q < 0, \quad S(i_d) = \lambda_m + \Delta i_d. \quad (8)$$

On $\mathcal{I}_{\text{op}}$, $\Delta i_d \geq 0$ and hence $S(i_d) \geq \lambda_m > 0$. The torque equality is therefore equivalent to

$$i_q(i_d) = \frac{C}{S(i_d)}. \quad (9)$$

Equation (9) is exact under the adopted model and reduces (6) to one scalar variable.

Equation (9) is the established torque-elimination coordinate used for the scalar LMC formulation in [20] and in other one-dimensional (1D) PMSM optimizations. Here it serves as the common coordinate transformation for every registered performance index, while the feasibility and solver safeguards remain independent of the selected objective.

For the numerical labels used later, T_a denotes the maximum torque attainable under the current constraint alone, and ω_a denotes the corresponding $R_s = 0$ voltage corner speed at that current-limited MTPA point. They serve only as normalization and scheduling bases; their values are specified with the test models in Section 8.

3. Convex Feasibility Reduction

3.1. Current Boundary

Substitution of (9) into (4) gives the current residual

$$F_I(i_d) = i_d^2 + \frac{C^2}{S(i_d)^2} - I_M^2. \quad (10)$$

Its derivatives are

$$F_I'(i_d) = 2i_d - \frac{2\Delta C^2}{S(i_d)^3}, \quad (11)$$

$$F_I''(i_d) = 2 + \frac{6\Delta^2 C^2}{S(i_d)^4} > 0. \quad (12)$$

Thus, the current-feasible sublevel set is an interval.

3.2. Full Voltage Boundary

Using (1), (2), and (9), the full voltage residual is

$$F_{V,R}(i_d) = R_s^2 i_d^2 + \omega^2 (\lambda_m + L_d i_d)^2 + \frac{(R_s^2 + \omega^2 L_q^2) C^2}{S(i_d)^2} + 2R_s \omega C - V_M^2. \quad (13)$$

Equivalently,

$$F_{V,R}(i_d) = A_R i_d^2 + B_R i_d + E_R + \frac{Q_R}{S(i_d)^2}, \quad (14)$$

where

$$A_R = R_s^2 + \omega^2 L_d^2, \quad B_R = 2\omega^2 \lambda_m L_d, \quad (15)$$

$$E_R = \omega^2 \lambda_m^2 + 2R_s \omega C - V_M^2, \quad Q_R = (R_s^2 + \omega^2 L_q^2) C^2. \quad (16)$$

The derivatives are

$$F'_{V,R}(i_d) = 2A_R i_d + B_R - \frac{2\Delta Q_R}{S(i_d)^3}, \quad (17)$$

$$F''_{V,R}(i_d) = 2A_R + \frac{6\Delta^2 Q_R}{S(i_d)^4} \geq 0. \quad (18)$$

The curvature is strictly positive except in degenerate zero-coefficient cases.

3.3. Zero-Resistance Voltage Boundary Without a Zero-Speed Singularity

When $R_s = 0$, the voltage residual can be written as

$$F_{V,0}(i_d) = \omega^2 \left[(\lambda_m + L_d i_d)^2 + \frac{L_q^2 C^2}{S(i_d)^2} \right] - V_M^2. \quad (19)$$

Unlike the flux-scaled form containing $(V_M/\omega)^2$, (19) remains defined at $\omega = 0$; it then reduces to the inactive constant residual $-V_M^2$. Its curvature is

$$F''_{V,0}(i_d) = 2\omega^2 L_d^2 + \frac{6\omega^2 L_q^2 \Delta^2 C^2}{S(i_d)^4} \geq 0. \quad (20)$$

3.4. Feasible-Interval Theorem

Proposition 1 (Connected scalar feasible set). *Under Assumption 1, define*

$$\mathcal{D}(T_r, \omega) = \{i_d \in \mathcal{I}_{\text{op}} : F_I(i_d) \leq 0, F_V(i_d) \leq 0\}, \quad (21)$$

where F_V is either (13) or (19). Then $\mathcal{D}(T_r, \omega)$ is a closed interval, a singleton, or the empty set.

Proof. The functions F_I and F_V are continuous and convex on $\mathcal{I}_{\text{op}}$. A scalar convex-function sublevel set is a convex subset of $\mathbb{R}$, hence an interval; continuity makes it closed. Intersecting the two sublevel intervals with the closed interval $\mathcal{I}_{\text{op}}$ preserves this property. $\square$

The existence of a scalar admissible domain was also used for the single LMC objective in [20]. Proposition 1 is retained because, in the present paper, the interval is the objective-independent interface reused by all registered indices; the new algorithmic issue is its certified construction and subsequent bracketed optimization.

When nonempty, write

$$\mathcal{D}(T_r, \omega) = [i_{d,\min}, i_{d,\max}]. \quad (22)$$

Figure 1 gives the geometric interpretation of Proposition 1. The detailed motor parameters of Model 1 considered in Figure 1 are given in Section 8. The current and voltage

sublevel sets retain only the highlighted segment of the torque-equality curve; mapping that segment through (9) gives all current pairs that satisfy the commanded torque and both limits. The constrained MTPA and mixed-index minimizers occupy different locations on this same segment, demonstrating that feasibility is objective independent whereas the selected performance index determines the issued reference. At this operating point, the left endpoint is imposed by the current limit and the right endpoint by the voltage limit. The mixed-index marker minimizes the normalized current–voltage stress objective introduced in Section 7.3, with the motor’s corner speed used as the scheduling transition.




Figure 1. Geometry for Model 1 at $\omega = 300$ rad/s and $T_r = 0.6T_a$ with $R_s = 0$. The green dotted curve is the torque-equality manifold, including infeasible portions. The thick red segment is the portion satisfying both current and voltage inequalities. Its left endpoint meets the blue current boundary and its right endpoint meets the orange voltage boundary. The purple circle and brown square are the constrained MTPA and mixed-index optima on this same feasible segment.

4. Current-First Feasibility and Positive-Torque Limiting

4.1. Voltage-Inactive Test

The current-feasible interval is first computed as

$$\mathcal{I}_C = \mathcal{I}_{\text{op}} \cap \{i_d : F_I(i_d) \leq 0\} = [i_{C,\min}, i_{C,\max}]. \quad (23)$$

If $\mathcal{I}_C = \emptyset$, the torque command is infeasible. Otherwise, because F_V is convex, the endpoint condition

$$F_V(i_{C,\min}) \leq 0, \quad F_V(i_{C,\max}) \leq 0 \quad (24)$$

implies $F_V(i_d) \leq 0$ for every $i_d \in \mathcal{I}_C$. The voltage stage is then skipped and $\mathcal{D} = \mathcal{I}_C$. If (24) fails, the voltage sublevel interval is computed and intersected with $\mathcal{I}_C$. This current-first sequence prevents irrelevant voltage roots outside the practical interval from clipping a low-speed solution.

4.2. Monotone Torque Feasibility

For $C \geq 0$, the residuals are nondecreasing with torque magnitude at every fixed i_d :

$$\frac{\partial F_I}{\partial C} = \frac{2C}{S(i_d)^2} \geq 0, \quad (25)$$

$$\frac{\partial F_{V,R}}{\partial C} = 2R_s\omega + \frac{2(R_s^2 + \omega^2 L_q^2)C}{S(i_d)^2} \geq 0, \quad (26)$$

$$\frac{\partial F_{V,0}}{\partial C} = \frac{2\omega^2 L_q^2 C}{S(i_d)^2} \geq 0. \quad (27)$$

Proposition 2 (Monotone positive-torque feasibility). *For fixed $\omega \geq 0$, if $0 \leq T_1 \leq T_2$, then $\mathcal{D}(T_2, \omega) \subseteq \mathcal{D}(T_1, \omega)$. Consequently, the feasibility indicator is monotone on the positive-torque axis.*

Proof. Equations (25)–(27) show that increasing $C = T/K_t$ cannot decrease either residual at any fixed i_d . Every point feasible at T_2 is therefore feasible at T_1 . $\square$

Define

$$\text{Feas}(T, \omega) = \begin{cases} 1, & \mathcal{D}(T, \omega) \neq \emptyset, \\ 0, & \mathcal{D}(T, \omega) = \emptyset. \end{cases} \quad (28)$$

A necessary precheck is $\text{Feas}(0, \omega) = 1$. If zero torque is infeasible, no positive torque can restore feasibility under the adopted branch; torque derating alone is insufficient. Otherwise, torque can be restricted to a finite interval by the current limit. The joint feasible set in this torque interval and the compact current interval is closed, and its projection onto the torque axis is therefore compact. Hence,

$$T_{\max}(\omega) = \max\{T \geq 0 : \text{Feas}(T, \omega) = 1\} \quad (29)$$

exists and can be found by scalar bisection. A conservative upper bracket is

$$T_{\text{hi}} = K_t I_M (\lambda_m + |\Delta| I_M), \quad (30)$$

and the effective command is

$$T_r^{\text{eff}} = \min\{T_r, \eta_T T_{\max}(\omega)\}, \quad 0 < \eta_T \leq 1. \quad (31)$$

For a finite-budget implementation that conservatively classifies near-tangent intervals, $\eta_T < 1$ prevents repeated evaluation exactly at the feasibility boundary; $\eta_T = 1$ is appropriate when the boundary is resolved by an exact or sufficiently certified offline calculation. An offline $T_{\max}(\omega)$ table can replace online torque bisection when a strict worst-case execution budget is required.

Figure 2 translates the monotone-feasibility result into the torque–speed envelope used by the outer limiter. The low-speed plateaus coincide with the current-only torque limit, whereas the descending portions begin when the voltage residual clips the feasible interval. Model 1 enters this voltage-limited regime earlier than Model 2 because its current-limited corner speed is lower; the plot therefore identifies the operating range in which torque derating, rather than objective selection alone, determines the command. The detailed motor parameters of Models 1 and 2 considered in Figure 2 are given in Section 8.




Figure 2. Maximum feasible positive torque for the two parameter sets using the $R_s = 0$ voltage model. The curves connect independently evaluated 5 rad/s speed samples as visual guides, whereas the markers denote independent evaluations at the four verification speeds reported in Section 8.2.

5. Safeguarded Boundary Computation

5.1. Taylor-Based Quadratic Initialization

The shared nonlinear term is

$$\frac{1}{S(i_d)^2} = \frac{1}{(\lambda_m + \Delta i_d)^2}. \quad (32)$$

A first-order expansion about $i_d = 0$ gives

$$\frac{1}{S(i_d)^2} \approx \frac{1}{\lambda_m^2} - \frac{2\Delta}{\lambda_m^3} i_d. \quad (33)$$

For the current boundary, (33) yields

$$i_d^2 - \frac{2\Delta C^2}{\lambda_m^3} i_d + \frac{C^2}{\lambda_m^2} - I_M^2 = 0. \quad (34)$$

With $\mu_I = \Delta C^2 / \lambda_m^3$, the two initial candidates are

$$i_{dI,0}^{\pm} = \mu_I \pm \sqrt{I_M^2 - \frac{C^2}{\lambda_m^2} + \mu_I^2}. \quad (35)$$

For a generic boundary

$$F(i_d) = A i_d^2 + B i_d + E + \frac{Q}{S(i_d)^2}, \quad (36)$$

the approximation gives

$$\tilde{B} = B - \frac{2\Delta Q}{\lambda_m^3}, \quad \tilde{E} = E + \frac{Q}{\lambda_m^2}, \quad (37)$$

and, when $A > 0$ and the discriminant is nonnegative,

$$i_{dV,0}^{\pm} = \frac{-\tilde{B} \pm \sqrt{\tilde{B}^2 - 4A\tilde{E}}}{2A}. \quad (38)$$

At $\omega = 0$ under the $R_s = 0$ model, the voltage endpoint test identifies the constraint as inactive, so (38) is not invoked with $A = 0$. Here, A , B , E , and Q are the scalar coefficients of the active boundary residual. For the current residual F_I , $(A, B, E, Q) = (1, 0, -I_M^2, C^2)$. For $F_{V,R}$, they are (A_R, B_R, E_R, Q_R) defined in (15) and (16); for $F_{V,0}$, they are $(\omega^2 L_d^2, 2\omega^2 \lambda_m L_d, \omega^2 \lambda_m^2 - V_M^2, \omega^2 L_q^2 C^2)$. The tildes in $\tilde{B}$ and $\tilde{E}$ denote the coefficients obtained after substituting the first-order approximation (33).

5.2. Candidate Refinement and Acceptance Test

Each lower or upper Taylor candidate from (35) or (38) is used as $i_d^{(0)}$ and refined independently. Here, k is the Newton-update index and N_B is the prescribed number of Newton updates applied to each boundary candidate:

$$i_d^{(k+1)} = i_d^{(k)} - \frac{F(i_d^{(k)})}{F'(i_d^{(k)})}, \quad k = 0, \dots, N_B - 1. \quad (39)$$

Fixed-count Newton refinement alone does not guarantee that the iterates remain on the desired root branch. Let $[a, b]$ denote the active candidate-certification interval: I_{op} for the current residual and I_C for the voltage residual. Define

$$s_F = \max\{F_{\text{floor}}, |F(a)|, |F(b)|\}. \quad (40)$$

Here, $F_{\text{floor}} = 1 \text{ A}^2$ for the current residual F_I , whereas $F_{\text{floor}} = 1 \text{ V}^2$ for the voltage residuals $F_{V,R}$ and $F_{V,0}$. After refinement, the two candidates are ordered as $i_{d,L} \leq i_{d,U}$, where the subscripts L and U denote the lower and upper boundary candidates, respectively. The index $j \in \{L, U\}$ selects either candidate. The tolerance ε_B , expressed in amperes, limits the magnitude of the residual-equivalent Newton correction used for boundary certification. The refined pair is accepted only when it satisfies all applicable conditions:

- (i) Ordered finite candidates within $[a, b]$.
- (ii) A midpoint residual $F((i_{d,L} + i_{d,U})/2) \leq 10^{-10} s_F$.
- (iii) The expected derivative orientation, $F'(i_{d,L}) < 0 < F'(i_{d,U})$, except when a boundary coincides with an already feasible search endpoint.
- (iv) For each non-endpoint boundary, the scaled residual and Newton-correction checks

$$F(i_{d,j}) \leq 10^{-12} s_F, \quad \left| \frac{F(i_{d,j})}{F'(i_{d,j})} \right| \leq \varepsilon_B, \quad j \in \{L, U\}. \quad (41)$$

A boundary coinciding with a search endpoint must satisfy $F \leq 10^{-10} s_F$.

The numerical study uses the boundary-location tolerance $\varepsilon_B = 10^{-3} I_M$. This is 0.01 A for Model 1 and 0.006 A for Model 2. It is a first-order boundary-location acceptance indicator, not the stopping tolerance of the objective solver or a model-uncertainty margin. The 10^{-10} midpoint/endpoint allowance and the stricter 10^{-12} root-residual allowance separate approximate candidate screening from the subsequent inward feasibility correction. They are relative to s_F , not absolute tolerances with interchangeable units.

Before any inward correction, the midpoint of the candidate joint interval must satisfy both residual inequalities with nonpositive evaluated residuals. If this feasible anchor is not available, the interval is reported as numerically unresolved and passed to the derating/status logic; a residual bisection must never start with two infeasible endpoints. Otherwise, each positive endpoint residual is corrected inward, using at most four bounded corrections followed, when needed, by N_R bisections toward that anchor. Floating-point equality noise remains possible when the physical quantities are recomputed in a different

arithmetic order. A candidate that fails a check is not issued. The implementation uses SI units and double precision. For other machine scales, a suitable representation is $z = i_d / I_M$, $f_I = F_I / I_M^2$, and $f_V = F_V / V_M^2$; derivatives must then include the corresponding chain-rule factors. A change of residual units must also rescale F_{floor} and dimensional derivative guards. The reference implementation guards a boundary Newton denominator below 10^{-16} and a candidate slope below 10^{-14} in its SI derivative units; these are numerical zero guards, not physical specifications. The objective curvature guard is 10^{-14} in the adopted objective units. Single precision or fixed-point arithmetic requires a new range/round-off analysis and retuned guards; the present double-precision tolerances cannot be transferred unchanged. The reference tolerance in Section 8.1 is used only for offline error evaluation. Section 8.12 reports direct sensitivity tests of the acceptance constants.

5.3. Deterministic Convex-Sublevel Fallback

For a rejected candidate, the fallback exploits convexity rather than reverting to a quartic formula. First, N_D derivative-bisection steps locate a fixed-budget approximation to the minimizer of F . If the resulting sampled minimum is positive, the sublevel set is conservatively classified as empty; this rule avoids returning a false feasible interval near tangency. Otherwise, each non-endpoint boundary is bracketed between an infeasible point and a feasible point and refined by N_R residual-bisection steps. The returned interval is checked for a joint feasible anchor and moved inward when necessary, subject to the finite-precision qualification above. A small torque margin $\eta_T < 1$ can be used when operation exactly at a tangential feasibility boundary is not required.

The fast path and fallback have fixed evaluation budgets. Here, safeguarded denotes candidate certification and deterministic fallback before the feasible interval is issued. Plant-uncertainty robustness requires separate modeling and validation.

For a simple boundary root r , suppose $|F'| \geq m_F > 0$ and $|F''| \leq M_F$ in a neighborhood containing the Newton iterates and r . Taylor's theorem gives

$$|x_{k+1} - r| \leq \frac{M_F}{2m_F} |x_k - r|^2.$$

Thus, Newton refinement is locally quadratic, but its unrestricted iterates do not have a global branch guarantee. At tangency, $F'(r) = 0$, so this local estimate cannot justify the fast path. For an established sign-changing bracket of width W , N_R bisections reduce the width to $W2^{-N_R}$ and return its feasible side. This linear contraction is independent of the Taylor initial guess.

The derivative bisection locating an interior residual minimum has location error at most $W2^{-(N_D+1)}$ in exact arithmetic. If $F'' \leq M_F$ on that bracket, the sampled minimum exceeds the true minimum by at most $M_F W^2 / 2^{2N_D+3}$. Therefore a positive sampled minimum can be unresolved at a near-tangent feasible point; its sign alone does not prove exact infeasibility. The implementation conservatively declines such a point within its fixed budget. These estimates explain both the suitability of bounded bisection as a recovery path and the need for supervisory derating near the feasibility boundary.

6. Convex Performance-Index Optimization

6.1. Reduced Objective and Boundary Test

For any performance index $J(i_d, i_q)$, define

$$\bar{J}(i_d) = J\left(i_d, \frac{C}{S(i_d)}\right), \quad i_d \in [i_{d,\min}, i_{d,\max}]. \quad (42)$$

Suppose $\bar{J}$ is strictly convex on the feasible interval. Let $g = \bar{J}'$ and $h = \bar{J}''$. The minimizer is the lower endpoint if $g(i_{d,\min}) \geq 0$, the upper endpoint if $g(i_{d,\max}) \leq 0$, and otherwise the unique zero of g lies in the interior.

6.2. Bracketed Newton Proposals with Mandatory Contraction

For an interior optimum, initialize the derivative-sign bracket $[\ell, u] = [i_{d,\min}, i_{d,\max}]$, where $g(\ell) < 0 < g(u)$. An objective-specific initial value is clipped to the bracket. Up to N_N Newton proposals

$$x_N = x - \frac{g(x)}{h(x)} \quad (43)$$

are evaluated only when $h(x) > 0$ and x_N lies strictly inside the current bracket; every evaluated point updates the bracket according to the sign of g . The algorithm then performs N_G mandatory midpoint derivative-bisection contractions and returns the midpoint. Here, N_N is the maximum number of objective-Newton proposals, N_G is the prescribed number of mandatory derivative-bisection contractions, x is the current scalar objective-search point, and x_N is its Newton trial. A finite trial strictly inside the current derivative-sign bracket is evaluated and used to contract the bracket; otherwise, the next search point is the current bracket midpoint.

Proposition 3 (Deterministic objective-bracket bound). *If the initial bracket contains the unique zero of the increasing derivative g , the safeguarded Newton updates cannot remove that zero. After N_G mandatory bisection contractions, the returned midpoint $\hat{i}_d$ satisfies*

$$|\hat{i}_d - i_d^*| \leq \frac{i_{d,\max} - i_{d,\min}}{2^{N_G+1}}. \quad (44)$$

where i_d^* is the unique exact minimizer of $\bar{J}$ on the feasible interval; in the interior case, it is equivalently the unique zero of g .

Proof. Every accepted point updates the bracket by the sign of a monotone function and therefore preserves the root. Each mandatory midpoint update halves the bracket width. The root-to-midpoint distance is no larger than half the final width. $\square$

The same local Newton argument applies to $g = \bar{J}'$ when $\bar{J}'' \geq m_J > 0$ and $|\bar{J}^{(3)}| \leq M_J$, giving local quadratic convergence with factor $M_J / (2m_J)$. The global search guarantee comes from preserving the derivative bracket and performing its mandatory bisections, regardless of whether a Newton proposal succeeds. Thus, the count can be selected before execution: for an initial width W and desired midpoint error $\epsilon_x > 0$, it suffices that $N_G \geq \max\{0, \lceil \log_2(W / (2\epsilon_x)) \rceil\}$. For the initial operating widths 10 A and 6 A, the conservative $N_G = 16$ gives bounds 7.63×10^{-5} A and 4.58×10^{-5} A before any benefit from Newton contraction. Proposition 3 assumes a valid bracket and correct derivative signs in exact arithmetic. When the implemented interval has approximate endpoints, its objective-search bound concerns the minimizer on that computed interval. The error relative to the exact model-constrained minimizer also includes boundary location error. If the two interval endpoints each differ by at most δ_{end} , strict convexity and the endpoint clipping rule give the bound $\delta_{\text{end}} + W2^{-(N_G+1)}$. The finite-precision tests do not provide a formal enclosure of δ_{end} ; it is assessed numerically in the reported reference comparisons. Newton convergence is not a closed-loop stability result.

Finally,

$$\hat{i}_q = \frac{C}{S(\hat{i}_d)}. \quad (45)$$

6.3. Objective-Registration Rule

The deterministic derivative-bracket path is used only when strict convexity, or an equivalent certificate of a unique minimizer and a strictly increasing derivative, is established on the actual feasible interval. Sampling the curvature at a few points is a diagnostic, not a proof. A lookup-table loss map, empirically fitted surface, or objective that fails an analytical convexity condition is routed to a direct bounded-domain scalar fallback, such as a prescribed grid scan followed by local bracket refinement. For a nonconvex or lookup-table objective, global optimality is claimed only when the registered fallback provides an appropriate global certificate over the complete feasible interval. Derivative bisection alone is reserved for objectives whose derivative monotonicity is established.

6.4. Overall Safeguarded Procedure

Algorithm 1 converts the preceding feasibility and objective results into one executable workflow. Its inputs are the torque and speed commands, motor data, practical limits, a registered objective, and fixed evaluation counts; its outputs are a feasibility status, the effective torque after any derating, and the current pair. The remaining count N_T is the number of outer torque-bisection contractions. The count tuple therefore distinguishes boundary Newton updates (N_B), residual-minimum contractions (N_D), boundary-root contractions (N_R), objective-Newton proposals (N_N), mandatory objective-bracket contractions (N_G), and torque-limit contractions (N_T). The procedure first constructs and certifies the feasible interval, invokes torque bisection only when that interval is empty, and then minimizes the selected objective on the nonempty interval. Taylor/Newton formulas generate low-cost candidates, while the original residuals, derivative orientations, and bracket tests determine whether those candidates are accepted.

Algorithm 1 Safeguarded positive-motoring current-reference generation

Input: $T_r \geq 0$, $\omega \geq 0$, motor parameters and limits, $\mathcal{I}_{op}$, a registered objective J , fixed counts $(N_B, N_D, N_R, N_N, N_G, N_T)$, boundary tolerance ε_B , and torque margin η_T .

Output: feasibility/derating status, effective torque T_r^{eff} , $(\hat{i}_d, \hat{i}_q)$, and boundary/objective fallback indicators.

- 1: Set $T_r^{\text{eff}} = T_r$, $C = T_r^{\text{eff}}/K_t$, and form F_I and F_V .
 - 2: Construct $\mathcal{I}_C$ from F_I using the Taylor/Newton candidates only if they pass (40) and (41); otherwise use the convex-sublevel fallback.
 - 3: If $\mathcal{I}_C = \emptyset$, set $\mathcal{D} = \emptyset$. Otherwise, apply (24). If voltage is inactive, set $\mathcal{D} = \mathcal{I}_C$; if voltage is active, construct its certified sublevel interval and intersect it with $\mathcal{I}_C$ to obtain $\mathcal{D}$.
 - 4: Validate the joint feasible anchor and apply bounded inward endpoint correction. If interval construction or this check fails, evaluate the zero-torque precheck with the same finite-budget procedure. If it fails, return a numerically unresolved branch status without issuing a current pair. Otherwise perform N_T torque-bisection steps (or query an offline map), retain the lower feasible torque estimate $\hat{T}_{\max}$, and set $T_r^{\text{eff}} = \min\{T_r, \eta_T \hat{T}_{\max}\}$. Recompute $C = T_r^{\text{eff}}/K_t$ and repeat the interval construction, joint-anchor check, and bounded inward endpoint correction once for that effective torque, without resetting it to T_r . If this final construction remains unresolved, return that status; do not enter an unbounded retry loop.
 - 5: Form $\bar{J}(i_d) = J(i_d, C/S(i_d))$ on the nonempty interval $\mathcal{D}$.
 - 6: If strict convexity, or an equivalent unique-minimizer certificate, is established on $\mathcal{D}$, apply the endpoint derivative test. For an interior optimum, preserve the derivative-sign bracket during at most N_N Newton proposals and then perform N_G mandatory derivative-bisection contractions.
 - 7: If the objective is not certified convex, use the registered bounded-domain scalar fallback over $\mathcal{D}$, without claiming global optimality unless the fallback supplies a valid global certificate.
 - 8: Enforce the final feasible-range limiter, set $\hat{i}_q = C/S(\hat{i}_d)$, and return the result with the effective torque and fallback flags.
-

An uncertified boundary candidate activates the convex-sublevel fallback and is not interpreted as physical infeasibility. The returned status distinguishes a feasible command, a torque-derated command, and an infeasible-branch status under the configured finite

evaluation budget. Because the fallback may conservatively classify a near-tangent sublevel set as empty, the latter status should not be interpreted as an exact model-level or plant-level infeasibility certificate unless the boundary has been resolved by a sufficiently accurate offline calculation. An excessive positive torque command is replaced by T_r^{eff} through the monotone torque limiter. The mathematical predicate $\widehat{\text{Feas}}$ in (28) refers to exact model feasibility. Its finite-budget implementation is denoted $\widehat{\text{Feas}}$: a false result means that a feasible interval was not established within that budget. Near tangency, the exact nesting theorem does not eliminate numerical classification uncertainty. The online lower feasible torque estimate should therefore be validated after derating; it is not a formally certified maximum-torque enclosure.

Every returned current is limited to $\mathcal{D}$ and then satisfies the torque equality by recomputing $\hat{i}_q = C/S(\hat{i}_d)$. The low-cost and conservative count tuples in Section 8 instantiate Algorithm 1; that section reports the returned-current accuracy, fallback use, and fixed evaluation structure.

6.5. Implementation Margins

When modeling error, current ripple, dc-link ripple, or inverter nonlinearity must be accommodated, the same residuals can use designer-selected effective limits

$$I_M^{\text{eff}} = \eta_I I_M, \quad V_M^{\text{eff}} = \eta_V V_M, \quad 0 < \eta_I, \eta_V \leq 1. \quad (46)$$

Alternatively, the final scalar interval can be contracted inward before objective minimization. The convexity and monotonicity results are unchanged after substituting the effective limits. No universal numerical margin is prescribed here because the required allowance depends on the inverter, sensing, ripple, thermal, and parameter-uncertainty budgets of the target drive.

7. Performance-Index Examples

7.1. Current Minimization and the Active-Power Index

A normalized MTPA index is

$$J_I = \frac{i_d^2 + i_q^2}{I_M^2}, \quad \bar{J}_I(i_d) = \frac{1}{I_M^2} \left(i_d^2 + \frac{C^2}{S(i_d)^2} \right). \quad (47)$$

Its curvature is

$$\bar{J}_I''(i_d) = \frac{1}{I_M^2} \left(2 + \frac{6\Delta^2 C^2}{S(i_d)^4} \right) > 0. \quad (48)$$

A low-cost initial value is $i_{d,0}^{\text{MTPA}} = \Delta C^2 / \lambda_m^3$. Recent MTPA studies have addressed parameter variation through data-driven offline maps with online magnet-temperature compensation [12] and have used real-time numerical optimization with nonlinear flux-linkage models [13]. Those methods improve the MTPA trajectory itself under nonideal parameter behavior. Equation (47) instead serves as one registered convex objective inside the present common constraint-handling framework. Under the adopted dq normalization, the active-power index satisfies

$$J_P = v_d i_d + v_q i_q = R_s(i_d^2 + i_q^2) + \omega\{\lambda_m + \Delta i_d\}i_q = R_s(i_d^2 + i_q^2) + \omega C. \quad (49)$$

For fixed T_r and ω , minimizing J_P is therefore equivalent to current minimization when $R_s > 0$, while J_P is constant on the torque-equality manifold when $R_s = 0$. It is not an independent objective in this model. This result is distinct from voltage-magnitude minimization, which measures voltage utilization rather than active input power.

7.2. Voltage Minimization as an MTPV-Dual Problem

The normalized voltage index is

$$J_V = \frac{v_d^2 + v_q^2}{V_M^2}. \quad (50)$$

For a fixed torque, minimizing (50) minimizes the required voltage. This is the dual fixed-torque interpretation of classical MTPV operation, which maximizes torque at a prescribed voltage limit. The reduced voltage objective is convex under the same conditions as the voltage residual. Recent FW reviews provide broader classifications of calculation-based, voltage-feedback, and predictive implementations [9,10]; the present objective is only the steady-state optimization layer and does not replace their transition, regulation, or overmodulation mechanisms. Except for the degenerate case $R_s = \omega = 0$, the reduced voltage objective is strictly convex under the stated model conditions. In the degenerate case, J_V is constant on the feasible interval, and a deterministic tie-breaking rule or another registered objective must be used.

7.3. Scheduled Current–Voltage Stress Index

A dimensionless mixed index is

$$J_\rho = (1 - \rho) \frac{i_d^2 + i_q^2}{I_M^2} + \rho \frac{v_d^2 + v_q^2}{V_M^2}, \quad 0 \leq \rho \leq 1, \quad (51)$$

with, for example,

$$\rho(\omega) = \frac{\omega^2}{\omega^2 + \omega_b^2}. \quad (52)$$

Here, $\omega_b > 0$ is the scheduling transition speed at which $\rho(\omega_b) = 1/2$; Section 8 sets $\omega_b = \omega_a$ for each evaluated motor. The two terms are squared utilization measures, with $I^2 = i_d^2 + i_q^2$ and $V^2 = v_d^2 + v_q^2$: I^2/I_M^2 penalizes current loading and, for a fixed R_s , is proportional to normalized copper loss; V^2/V_M^2 penalizes the required steady-state inverter voltage and favors voltage headroom. The voltage term is not electrical power, switching loss, or a measured insulation or thermal-aging law. Their weighted sum is a design index for choosing among torque-equivalent feasible currents. Current and voltage remain hard constraints for every weight, so a lower index cannot authorize a limit violation. The choice $\omega_b = \omega_a$ uses the motor's available characteristic speed as a reproducible transition scale: $\rho(0) = 0$, $\rho(\omega_a) = 1/2$, and $\rho \rightarrow 1$ at high speed. It favors current economy at low speed and increases the importance of voltage headroom as the back-EMF demand grows. It is a heuristic schedule, not a derived optimum for a loss model and not the exact full- R_s field-weakening transition. In particular, ω_a stays nominal in the voltage and parameter-offset tests; a controller may instead schedule the weight using measured voltage utilization, but that would be a different design requiring validation. Section 8.10 shows the resulting current/voltage trade-off numerically.

Let $a_\rho = (1 - \rho)/I_M^2$ and $b_\rho = \rho/V_M^2$. Using the coefficients in (14), a Taylor stationarity estimate is

$$i_{d,0}^\rho = \frac{\Delta(a_\rho C^2 + b_\rho Q_R)/\lambda_m^3 - b_\rho B_R/2}{a_\rho + b_\rho A_R}. \quad (53)$$

For the $R_s = 0$ model, A_R , B_R , and Q_R are replaced by $\omega^2 L_d^2$, $2\omega^2 \lambda_m L_d$, and $\omega^2 L_q^2 C^2$, respectively.

7.4. Loss-Oriented Objectives

If $J = P_{\text{loss}}(i_d, i_q, \omega)$, the framework becomes a constrained LMC layer. Here, P_{loss} denotes the selected steady-state analytical or fitted loss function; it is not assumed to be a measured total-drive loss unless separately identified and validated. Recent LMC and efficiency-oriented studies include analytical current allocation with iron-loss resistance [14], bivariate-extreme-value optimization [15], linearized torque/loss modeling [16], and equivalent-circuit-based drive-system design optimization [17]; harmonic-current design can additionally trade torque ripple against iron loss [22]. The residual-corrected method in [20] is the closest single-objective scalar antecedent, but it uses one adopted loss model and one correction layer. In the present framework, an analytical copper-plus-iron-loss model can use the convex path only when its reduced curvature is proven positive. An experimentally fitted loss map is not assumed convex and uses the direct scalar fallback unless a valid interval certificate is available. Model-predicted loss reduction supports comparison of the adopted objective; drive-efficiency improvement requires corresponding measurements. This is only an optimization layer: even exact minimization of an inaccurate loss model may produce poor efficiency choices in the machine. Final control performance depends on the accuracy and operating-range coverage of the adopted loss model as well as current tracking, voltage headroom, and supervisory transitions.

7.5. Structurally Convex Design Extensions

To avoid the dimensional ambiguity of an unnormalized total-loss index, define a thermal-priority current–flux surrogate

$$\bar{J}_{\text{th}}(i_d) = w_c \frac{i_d^2 + C^2/S(i_d)^2}{I_M^2} + w_f(\omega) \frac{(\lambda_m + L_d i_d)^2 + L_q^2 C^2/S(i_d)^2}{\Psi_B^2}, \quad (54)$$

where $w_c > 0$ and $w_f(\omega) \geq 0$ are dimensionless design weights, and $\Psi_B > 0$ is a flux-linkage normalization base expressed in webers. Its curvature is

$$\bar{J}_{\text{th}}'' = \frac{w_c}{I_M^2} \left(2 + \frac{6\Delta^2 C^2}{S(i_d)^4} \right) + \frac{w_f}{\Psi_B^2} \left(2L_d^2 + \frac{6L_q^2 \Delta^2 C^2}{S(i_d)^4} \right) > 0. \quad (55)$$

The index can encode a design priority, but it is not called a measured thermal loss unless the weights are identified from a thermal/loss model.

A normalized d -axis flux-target penalty is

$$\bar{J}_{\psi}(i_d) = \frac{i_d^2 + C^2/S(i_d)^2}{I_M^2} + \beta \left(\frac{\lambda_m + L_d i_d - \psi_d^{\text{tgt}}}{\Psi_B} \right)^2, \quad \beta > 0, \quad (56)$$

with curvature

$$\bar{J}_{\psi}'' = \frac{1}{I_M^2} \left(2 + \frac{6\Delta^2 C^2}{S(i_d)^4} \right) + \frac{2\beta L_d^2}{\Psi_B^2} > 0. \quad (57)$$

Both indices are structurally convex and use the same feasibility layer. Here, ψ_d^{tgt} is the designer-selected target d -axis flux linkage [Wb], and $\beta > 0$ is a dimensionless penalty weight.

7.6. Conditional Parameter-Error Objective

For independent small errors in λ_m and Δ , a torque-error variance proxy is

$$J_{\text{par}} = i_q^2(\sigma_{\lambda}^2 + \sigma_{\Delta}^2 i_d^2), \quad \bar{J}_{\text{par}} = \frac{C^2(\sigma_{\lambda}^2 + \sigma_{\Delta}^2 i_d^2)}{S(i_d)^2}. \quad (58)$$

Its curvature is

$$\bar{J}_{\text{par}}''(i_d) = \frac{C^2}{S(i_d)^4} \left[6\sigma_\lambda^2 \Delta^2 + 2\sigma_\Delta^2 \lambda_m (\lambda_m - 2\Delta i_d) \right]. \quad (59)$$

Here, σ_λ [Wb] and σ_Δ [H] are the prescribed standard deviations of the independent small errors in λ_m and Δ , respectively. Because the bracketed term increases with i_d for $\Delta < 0$, its minimum occurs at $i_{d,\min}$. For $C > 0$, (58) is strictly convex on the entire feasible interval when the bracketed term in (59) is positive at $i_{d,\min}$ and convex when it is nonnegative there. At $C = 0$, $\bar{J}_{\text{par}} = 0$, so the objective has no unique minimizer and is not admitted to the derivative-bracket fast path unless a deterministic tie-breaking rule is registered. This analytical endpoint test replaces curvature sampling as the admission criterion. The objective is a local variance proxy. Parameter observers, saturation maps, and robust predictive control address complementary system-level functions [12,23].

Figure 3 visualizes the resulting objective-registration decision. In Figure 3a, the three curvature curves remain positive over the complete feasible interval, so the endpoint derivative test and bracketed fast path are valid. In Figure 3b, increasing σ_Δ drives the parameter-error curvature below zero over part of the interval; the scalar reduction remains valid, but the convex fast path must be rejected in favor of the registered bounded-domain scalar fallback. The figure therefore distinguishes objective compatibility with the 1D framework from eligibility for its faster convex solver.




Figure 3. Reduced-objective curvature for Model 1 at $\omega = 300$ rad/s and $T_r = 0.6T_a$. Panel (a) uses the thermal-priority and flux-target settings reported in Table 1. Panel (b) uses $\sigma_\lambda = 0.02\lambda_m$ and the indicated σ_Δ levels. The dashed horizontal line marks zero curvature; a curve below this line remains compatible with scalar reduction but is routed away from the convex fast path.

Table 1 quantifies the minimum curvature, rather than a visually selected point, over five representative feasible intervals. Every reported minimum is positive, which certifies the convex fast path for the normalized MTPA, thermal-priority, and flux-target examples at those operating points. The different magnitudes indicate different local conditioning, but they do not change the uniqueness of the interval minimizer.

Table 1. Minimum reduced-objective curvatures [A^{-2}] for the structurally convex examples. The thermal-priority example uses $w_c = 1$, $w_f = [\omega/(300 \text{ rad/s})]^2$, and $\Psi_B = \lambda_m$; the flux-target example uses $\beta = 5$ and $\Psi_B = \lambda_m$.

Operating Point	Feasible Interval [A]	Normalized MTPA	Thermal Priority	Flux Target
Model 1, $\omega = 100, 0.3T_a$	$[-9.841, 0.000]$	2.0865×10^{-2}	5.7755×10^{-2}	1.5302
Model 1, $\omega = 300, 0.6T_a$	$[-9.279, -4.383]$	2.4041×10^{-2}	4.6666×10^{-1}	1.5333
Model 1, $\omega = 400, 0.3T_a$	$[-8.412, -1.261]$	2.1299×10^{-2}	6.3838×10^{-1}	1.5306
Model 2, $\omega = 100, 0.9T_a$	$[-4.573, -1.548]$	8.5469×10^{-2}	1.3901×10^{-1}	1.0200
Model 2, $\omega = 400, 0.9T_a$	$[-4.573, -2.734]$	8.5469×10^{-2}	9.4204×10^{-1}	1.0200

Table 2 applies the analytical endpoint condition to the conditional parameter-error objective. A $2\%|\Delta|$ saliency-error weight leaves only a marginally positive minimum curvature, whereas weights of $5\%|\Delta|$ and above make the minimum negative. These entries define an explicit solver-routing threshold: the former case can use the convex path with care, while the latter cases require the bounded-domain scalar fallback.

Table 2. Convexity of (58) at Model 1, $\omega = 300 \text{ rad/s}$, and $T_r = 0.6T_a$, with $\sigma_\lambda = 0.02\lambda_m$. Curvature values are in Wb^2 under the adopted variables.

σ_Δ	Minimum $\bar{J}_{\text{par}}''$	Convex Fast Path?
$2\% \Delta $	$+5.147 \times 10^{-7}$	Yes (marginal)
$5\% \Delta $	-3.416×10^{-5}	No
$10\% \Delta $	-1.896×10^{-4}	No
$15\% \Delta $	-4.551×10^{-4}	No

8. Numerical Verification

8.1. Models, Reference Solution, and Test Matrix

The selected objective is (51) with $\omega_b = \omega_a$ for each motor. The high-accuracy reference uses the same scalar convex structure only for offline error evaluation. For each active boundary residual, 120 derivative-bisection contractions locate the residual minimum and 120 residual-bisection contractions refine each required boundary root. After the endpoint derivative test, an interior objective minimizer is resolved by 130 derivative-bisection contractions. In this section, i_d^* denotes the high-accuracy numerical approximation to the exact minimizer defined in Proposition 3. All numerical tests used $I_{\text{op}} = [-I_M, 0]$; no additional demagnetization bound was imposed. The nominal current and voltage limits in Table 3 were used directly, corresponding to $\eta_I = \eta_V = 1$, and no additional inward engineering margin was applied apart from the numerical endpoint-feasibility correction described in Section 5.2. All reference calculations used double-precision arithmetic. For reference feasibility classification, the residual-minimum tolerance was $f_{\text{tol}} = \max\{10^{-13}, 5 \times 10^{-12} \max(1, |F(a)|, |F(b)|)\}$. A residual minimum exceeding this tolerance was classified as infeasible, whereas a near-tangent case within the tolerance was retained as a singleton feasible interval. The high-accuracy reference and the proposed method solve the same reduced, constant-parameter steady-state optimization problem. Accordingly, their current-reference and objective-value differences quantify numerical solution accuracy conditional on that model; they do not constitute independent validation of motor-drive dynamics, physical parameter accuracy, or measured efficiency. The parameter-offset study in Section 8.8 separately examines the consequences of using incorrect magnetic parameters and retains the same steady-state scope.

The two parameter sets in Table 3 are taken from [24]. The simplified $R_s = 0$ model is used for the representative comparison to match that source, and both (13) and (19) are

used in the broader sweep. The values of T_a and ω_a were used directly as the torque-normalization and scheduling constants in all numerical calculations.

Table 3. Motor parameters and normalization/scheduling constants used in the numerical verification [24].

Parameter	Model 1	Model 2
R_s [Ω]	2.48	1.10
L_d [mH]	74.98	59.00
L_q [mH]	113.91	101.00
λ_m [Wb]	0.193	0.193
I_M [A]	10	6
V_M [V]	196.7	196.7
T_a [N m]	10.2387	4.9856
ω_a [rad/s]	206.90	385.91
Number of poles p	4	4

Table 3 also shows why the two validation models are complementary rather than duplicate scalings. They share the same magnet flux linkage and voltage limit, but differ in current rating, inductances, saliency ratio, and corner speed; Model 1 reaches voltage-limited operation earlier, whereas Model 2 retains its current-limited plateau to a higher speed. The commanded fractions are normalized by T_a , and the mixed-index schedule uses $\omega_b = \omega_a$ in (52), so the test matrix compares the same normalized control task across these different electrical scales.

The representative matrix contains $\omega \in \{100, 200, 300, 400\}$ rad/s and $T_r \in \{0.3T_a, 0.6T_a, 0.9T_a\}$ for each motor. Of the 24 combinations, 21 are feasible. Model 1 is infeasible at $(0.6T_a, 400)$, $(0.9T_a, 300)$, and $(0.9T_a, 400)$ under the $R_s = 0$ voltage limit.

8.2. Torque-Limiter Check

Table 4 reports the model-level maximum feasible torque obtained by 55 outer torque-bisection contractions using the high-accuracy interval-existence predicate described in Section 8.1. The torque search uses $[0, T_a]$ as its initial bracket, and the reported value is the final lower feasible endpoint. Table 4 connects the monotonicity result in Proposition 2 to the actual test matrix.

Table 4. Maximum feasible positive torque [N m] obtained with 55 scalar torque-bisection steps and the $R_s = 0$ model.

Model	$\omega = 100$	$\omega = 200$	$\omega = 300$	$\omega = 400$
Model 1	10.238689	10.238689	6.935897	4.723750
Model 2	4.985562	4.985562	4.985562	4.969762

At 100 and 200 rad/s, both models attain their current-only torque base T_a , so the voltage constraint is inactive at the maximum feasible torque. For Model 1, $T_{\max}$ falls to 6.935897 N m at 300 rad/s and 4.723750 N m at 400 rad/s; these values exclude $0.9T_a$ at 300 rad/s and both $0.6T_a$ and $0.9T_a$ at 400 rad/s. Model 2 remains almost current limited through 400 rad/s. Table 4 therefore explains all three infeasible representative cases and verifies that the scalar bisection locates the model-consistent derating boundary.

8.3. Representative Accuracy and the Need for Safeguards

Two safeguarded settings are evaluated:

$$(N_B, N_D, N_R, N_N, N_G, N_T) = \begin{cases} (6, 12, 12, 3, 12, 55), & \text{low-cost,} \\ (8, 14, 14, 4, 16, 55), & \text{conservative.} \end{cases}$$

For comparison, an unchecked fixed-Newton baseline and two direct grid/bisection baselines are retained. Fallback cases in Table 5 count operating points at which at least one boundary candidate was rejected and the convex-sublevel fallback was used. For the direct grid/bisection baselines, N_{grid} is the number of uniform grid subintervals and N_{bis} is the number of residual-bisection contractions used to refine each detected feasible boundary. Let $\hat{i}_d$ be the value returned by an evaluated method and i_d^* the high-accuracy reference. The reported relative objective error is

$$e_J = 100 \frac{|\bar{J}(\hat{i}_d) - \bar{J}(i_d^*)|}{\max\{|\bar{J}(i_d^*)|, 10^{-15}\}} \%. \quad (15)$$

The denominator floor is a numerical guard for the dimensionless mixed objective.

Table 5. Accuracy over the 21 feasible representative cases relative to the high-accuracy scalar reference. (n/a: The direct grid/bisection baselines do not use the proposed boundary-fallback mechanism).

Method	Mean $ \hat{i}_d - i_d^* $ [A]	Max. $ \hat{i}_d - i_d^* $ [A]	Max. e_J Error	Boundary Fallback
Direct grid/bisection, $N_{\text{grid}} = 16, N_{\text{bis}} = 5, N_G = 8$	0.005408	0.013577	0.003192%	n/a
Direct grid/bisection, $N_{\text{grid}} = 32, N_{\text{bis}} = 6, N_G = 12$	0.000342	0.001174	$1.019 \times 10^{-5}\%$	n/a
Unchecked fixed Newton, $N_B = 6, N_N = 1$	0.512911	1.536693	12.381%	n/a
Unchecked fixed Newton, $N_B = 6, N_N = 3$	0.001273	0.015567	0.001093%	n/a
Safeguarded low-cost	0.000189	0.000510	$1.268 \times 10^{-6}\%$	10/21
Safeguarded conservative	0.000019	0.000046	$1.747 \times 10^{-8}\%$	3/21

Table 5 isolates the value of the safeguards. One unchecked objective update reaches a 1.536693 A maximum error, and three unchecked updates reduce that error but still provide no certificate. The safeguarded low-cost and conservative settings reduce the maximum errors to 0.000510 and 0.000046 A, respectively. Their fallback counts, 10/21 and 3/21, show that candidate certification actively rejects questionable boundary estimates and that additional fixed refinements reduce the need for fallback. The relevant gain is therefore not Newton iteration alone, but the combination of candidate checking, deterministic recovery, and the bracket bound in (44).

8.4. Broad Grid and Fixed-Seed Stress Test

For each motor and each voltage model, a 31×31 grid was evaluated. Model 1 used $\omega \in [25, 500]$ rad/s, Model 2 used $\omega \in [25, 700]$ rad/s, and both used $T_r/T_a \in [0.05, 1]$. The combined matrix contains 3844 points, of which 3071 are feasible according to the high-accuracy reference.

Table 6 aggregates current-reference error and fallback activation over all 3071 feasible points, testing whether the representative-case behavior persists across the full evaluated speed–torque domains. The value $N_T = 55$ applies only when the outer torque limiter is invoked. Tables 5 and 6 report current-reference errors only for commands classified as

feasible by the high-accuracy reference; therefore, the outer torque limiter, torque-command derating, and N_T do not enter those error statistics.

Table 6. Aggregate accuracy over the 3071 feasible points in the 3844-point systematic sweep.

Setting	Mean $ \hat{i}_d - i_d^* $ [A]	Max. $ \hat{i}_d - i_d^* $ [A]	Max. e_J Error	Fallback Cases	Fallback Rate
Low-cost	0.000192	0.001331	0.003708%	895	29.144%
Conservative	0.000016	0.000289	0.000341%	149	4.852%

Table 6 shows that the conservative setting lowers the maximum error from 0.001331 to 0.000289 A and the fallback rate from 29.144% to 4.852% by performing more fixed refinements before certification. Figure 4 resolves the conservative error by speed: Model 1 exhibits two localized peaks near the middle of its evaluated range, whereas Model 2 remains nearly flat, and the $R_s = 0$ and full- R_s curves remain close for both motors. The branch-level diagnosis in Section 8.9 associates the Model 1 peaks with voltage-boundary fallback and a boundary optimum; they are not unexplained drive-level errors. Every plotted value lies far below the $0.5\%I_M$ guide, so the broad sweep supports the fixed-budget accuracy claim throughout the evaluated grids rather than only at the 21 representative points. The dotted line is used only as an engineering-scale visual reference, not as a solver stopping tolerance or a universal design requirement.




Figure 4. Worst d -axis current-reference error at each speed for the conservative setting. At each speed, the ordinate is the maximum of $|\hat{i}_d - i_d^*|$ over the feasible cases among the 31 evaluated torque fractions $T_r/T_a \in [0.05, 1]$. The vertical axis is logarithmic, and the curves connect the 31 evaluated speed samples as visual guides. Solid and dashed curves correspond to the $R_s = 0$ and full- R_s voltage models, respectively. The dotted $0.5\%I_M$ line is an engineering comparison guide rather than a solver tolerance or a universal design requirement.

A separate fixed-seed randomized stress test used 5000 draws for each motor-voltage-model subset and each safeguarded setting. Speed was sampled from $[0, 650]$ rad/s for Model 1 and $[0, 850]$ rad/s for Model 2; T_r/T_a was sampled from $[0, 1.15]$. Each setting therefore used 20,000 draws. Because the same fixed seeds were reused for both solver settings, each setting contained the same 13,046 feasible reference cases. Neither setting produced a missed feasible case or a false feasible case. The maximum d -axis errors were 0.001259 A for the low-cost setting and 0.000310 A for the conservative setting. These results characterize the tested parameter boxes.

8.5. Evaluation-Count Structure

Table 7 separates the evaluation burden of interval construction, fallback, objective minimization, and torque limiting. This separation is needed because the manuscript claims a fixed evaluation structure, not measured execution time on a particular processor.

Table 7. Fixed-count evaluation structure. Counts report the core update/contraction evaluations; fixed endpoint and certification checks are listed separately.

Stage	Derivative Evaluations	Residual/Curvature Evaluations	Comment
Taylor/Newton boundary fast path, one active constraint	$2N_B$	$2N_B + \text{certification}$	Two candidate roots; voltage stage is skipped when inactive.
Convex-sublevel boundary fallback	N_D	up to $2N_R + \text{endpoint checks}$	Locates the convex minimum and refines the sublevel boundaries.
Safeguarded objective minimization	2+ up to $2N_N + N_G$ first derivatives	N_N second derivatives	N_G mandatory contractions give (44).
Outer torque limiter	0	N_T interval-existence tests	Optional online layer; an offline map removes this repeated cost.
Final joint-feasibility projection	up to 16	up to $30 + 4N_R$	Two joint-anchor residual checks; up to four corrections per endpoint, then, if needed, N_R joint-residual bisections.

The boundary and objective stages in Table 7 are completely determined by $(N_B, N_D, N_R, N_N, N_G)$, including the mandatory N_G contractions that produce (44). Only the online torque limiter repeats the interval-existence test N_T times, and an offline $T_{\max}(\omega)$ map removes that repeated layer. The table summarizes the fixed-count core and the bounded certification overhead of the reported solver settings. Together, Tables 5–7 trace Algorithm 1 from returned-current accuracy and fallback activation to its fixed evaluation budget.

8.6. Interpretation of the Results

The systematic and randomized accuracy results in this section evaluate the scheduled mixed current–voltage objective in (51). Section 8.10 separately verifies MTPA and voltage minimization numerically. The other design objectives in Section 7 demonstrate the registration and solver-routing structure. Grid/bisection provides a strong reference and fallback. The proposed structure avoids the initial grid scan for objectives admitted to the analytical path and makes the active constraint boundaries explicit. The representative results also show why an unchecked one-step Newton approximation is unsuitable for a general mixed objective: its maximum current error exceeded 1.5 A. Candidate certification, fallback, and a mandatory derivative bracket are therefore substantive parts of the method rather than implementation details.

The feasible-interval decomposition also clarifies objective reuse. MTPA, voltage minimization, and the two convex extensions select different points from the same interval, whereas the current/voltage feasibility calculation is unchanged. The parameter-error example demonstrates the boundary of this generality: the objective is still 1D, but it must use a registered bounded-domain scalar fallback when its analytical curvature condition fails. Global optimality is claimed only when that fallback supplies an appropriate certificate over the complete feasible interval.

The numerical safeguards address boundary construction and scalar minimization for the parameters supplied to the algorithm. They do not estimate magnetic parameters or correct discrepancies between that model and the physical machine. In particular, a small error relative to the high-accuracy scalar reference must not be interpreted as a small

torque-tracking error in a drive with uncertain parameters. Section 8.8 distinguishes these two error sources explicitly.

8.7. Relationship to Prior Methods

Table 8 compares the approaches by functional role rather than by a universal ranking. Its purpose is to identify where the present optimization layer fits within a complete drive and which functions—such as region management, parameter adaptation, or current-loop control—remain the responsibility of surrounding methods.

Table 8. Functional comparison of related current-reference approaches and their best-matched roles.

Approach	Main Strength	Scope or Limitation Relevant Here	Best-Matched Role
Analytical quartic/unified theory [6,7]	Exact analytical structure for specified machine models and objectives	Expressions and branch management can be elaborate; changing the objective requires corresponding analytical development	Offline analysis or specialized analytical law
Whole-region numerical torque control [8]	Covers MTPA/FW/maximum-current/MTPV regions and has prior experimental validation	Primarily designed around minimum-current torque control and complete drive-level region logic	Complete minimum-current torque-control scheme
Recent parameter-aware or numerical MTPA [12,13]	Accounts for temperature, saturation, cross-coupling, or nonlinear flux-linkage behavior	Targets MTPA accuracy rather than a reusable multi-objective feasibility/optimization interface	Accurate MTPA tracking under parameter variation
Recent efficiency/LMC methods [14–17]	Uses richer loss models, analytical allocation, or system-level optimization	Objective and implementation are loss specific; evidence applies to the corresponding calibrated loss model	Efficiency-oriented control or design
Explicit algebraic approximations [24]	Very low online arithmetic for selected MTPA, MTPV, and LMC objectives	Approximation formulas are objective specific	Fast specialized reference law
Residual-corrected scalar LMC [20]	Uses an exact scalar LMC residual to refine an explicit warm start at fixed correction cost	Specialized to one adopted LMC objective and one correction layer	Low-cost LMC refinement
SVS anchor scheduling [21]	Reacts to speed and voltage margin through a deterministic scheduled base map	Scheduled candidate generation with projection, distinct from performance-index minimization	Cost-sensitive wide-speed base reference
Direct scalar grid/bisection	Derivative free and applicable to nonconvex or lookup-table objectives	Cost depends on grid and tolerance; active-boundary structure is not exploited	Reference, fallback, or offline map
Present framework	Reuses a certified feasible interval for selectable analytically convex indices and retains a bounded-domain scalar fallback	Positive motoring and steady-state constant parameters; hardware timing not measured	Modular current-reference optimization layer

Table 8 places the framework as a modular optimization layer between specialized low-cost reference laws and complete whole-region controllers. It supplies a common feasible interval and objective solver; the surrounding drive provides current regulation, dc-link management, reference transitions, protection, and supervisory torque/speed control.

8.8. Constant-Parameter Offsets and Frozen-Model Mismatch

8.8.1. Test Design and Interpretation

We distinguish numerical sensitivity to changed model coefficients from sensitivity to an unmodeled plant change. Let θ_0 denote the nominal magnetic parameters and θ_a the parameters of the perturbed steady-state evaluation model. The offsets are

$$L_{d,a} = L_{d,0}(1 + \delta_d), \quad L_{q,a} = L_{q,0}(1 + \delta_q), \quad \lambda_{m,a} = \lambda_{m,0}(1 + \delta_\lambda). \quad (60)$$

At each amplitude $a \in \{0.10, 0.20\}$, six individual-parameter cases use one of $\delta_d, \delta_q, \delta_\lambda = \pm a$ with the other two zero; eight joint cases use all sign combinations in $\{-a, +a\}^3$. Including the nominal case gives 29 scenarios. These are deterministic sensitivity cases, not measured saturation characteristics or a probabilistic model. Every tested parameter set satisfies $L_{q,a} > L_{d,a} > 0$ and $\lambda_{m,a} > 0$.

For each motor and scenario, the full-resistance voltage model is evaluated on 31 uniformly spaced speeds (25–500 rad/s for Model 1 and 25–700 rad/s for Model 2) and 31 uniformly spaced torque fractions $T_r/T_a \in [0.05, 1]$. The values of R_s , I_M , V_M , p , T_a , and $\omega_b = \omega_a$ remain nominal. Thus, the commanded torques and the mixed-index schedule are identical across scenarios. Both count settings from Section 8.3 and the tolerances in Section 5.2 are retained. The test contains 55,738 motor/scenario/operating-point combinations for each setting. Torque derating is disabled in this fixed-command sensitivity test: infeasible commands are counted separately, and returned commands use $T_r^{\text{eff}} = T_r$.

The added calculations were executed in Python 3.13 using the archived safeguarded numerical routines and the 120/120/130 bisection reference described in Section 8.1. Before these tests, the nominal 3844-point sweep reproduced the reported 3071 feasible cases and 895/149 boundary-fallback counts. An additional comparison of 491 reference solutions with Brent root finding gave a maximum d -axis difference below 1.65×10^{-14} A. This cross-check concerns numerical root finding within the same physical formulation.

8.8.2. Numerical Accuracy When the Model Is Updated

First, both the solver and its high-accuracy reference use θ_a . The errors e_d and e_q are the absolute differences between their respective d - and q -axis current references at the same torque and speed. Table 9 pools the two motors and, at each nonzero amplitude, all 14 offset scenarios. Fallback rates use the reference-feasible count as the denominator. Missed and false feasible classifications are assessed over all tested commands, including infeasible ones.

Table 9. Updated-model numerical accuracy with the full-resistance voltage model. The nominal rows contain 1922 tested commands per setting; each 10% or 20% row contains 26,908. Errors are relative to the perturbed-model reference and are in amperes.

Offset	Setting	Feasible	Max. e_d [A]	Max. e_q [A]	Fallback [%]
Nominal	Low-cost	1523	1.3313×10^{-3}	6.7732×10^{-4}	28.9560
Nominal	Conservative	1523	2.2183×10^{-4}	1.0816×10^{-4}	4.7275
10%	Low-cost	20,870	1.4785×10^{-3}	6.7704×10^{-4}	29.3723
10%	Conservative	20,870	3.6737×10^{-4}	1.6718×10^{-4}	6.6267
20%	Low-cost	20,199	1.6138×10^{-3}	7.0987×10^{-4}	29.3430
20%	Conservative	20,199	4.0401×10^{-4}	1.7594×10^{-4}	8.7381

No missed or false feasible case was observed for either setting on this matrix. Recomputing i_q through the torque equality kept the modeled torque residual below 1.8×10^{-15} N m. Direct current/voltage checks showed only floating-point equality noise, with positive squared residuals no larger than 7.11×10^{-15} A² and 2.92×10^{-11} V². These findings support numerical accuracy for the tested constant coefficients. They do not establish robustness to unknown plant parameters. Aggregate fallback rates also conceal scenario variation: among individual motor/scenario subsets, the 20% cases reached 75.64% for the low-cost setting and 47.04% for the conservative setting. Such variation must be included when assessing computational cost.

8.8.3. Consequences of a Frozen Nominal Controller

Second, the issued currents $\hat{i}_0 = (\hat{i}_{d,0}, \hat{i}_{q,0})$ are computed with θ_0 , while torque and voltage are evaluated with θ_a . No parameter adaptation is applied. Let i_a^* be the high-accuracy mixed-index optimum for θ_a at the same command, speed, limits, and fixed schedule. For commands feasible under both parameter sets, define

$$e_{d,\text{mis}} = |\hat{i}_{d,0} - i_{d,a}^*|, \quad e_{q,\text{mis}} = |\hat{i}_{q,0} - i_{q,a}^*|. \quad (61)$$

These are current-reference discrepancies from a model-informed optimum, not simulated current-loop tracking errors. A command infeasible for the perturbed model has no such optimum and is recorded separately.

The steady-state torque-reference error is $e_T = T_{e,a}(\hat{i}_0) - T_r$. Under ideal current following, and using $\Delta_0 = L_{d,0} - L_{q,0}$ and $\Delta_a = L_{d,a} - L_{q,a}$, the torque equality gives the exact algebraic identity

$$\frac{e_T}{T_r} = \frac{(\lambda_{m,a} - \lambda_{m,0}) + (\Delta_a - \Delta_0)\hat{i}_{d,0}}{\lambda_{m,0} + \Delta_0\hat{i}_{d,0}}, \quad T_r > 0. \quad (62)$$

This expression is evaluated without a small-offset approximation. The required voltage magnitude $V_a = \sqrt{v_{d,a}^2 + v_{q,a}^2}$ is independently recomputed from (1)–(2) with θ_a . We report the normalized torque error $E_T = 100e_T/T_r$ in percent. Cases exceeding $V_M(1 + 10^{-10})$ are flagged; the small reporting allowance suppresses floating-point equality noise and does not alter the algorithm. The algebraic torque at a voltage-infeasible reference is not interpreted as dynamically achievable tracking.

Table 10 reports the conservative setting, separating individual offsets from simultaneous sign corners. Current errors use the common-feasible subset; torque and voltage statistics include all nominally issued references. The issued count minus the common count gives the number of commands infeasible for the perturbed model.

Table 10. Frozen nominal controller evaluated with perturbed magnetic parameters. Maxima are taken over the stated finite test matrix. “Common” counts commands feasible for both models; “issued” counts nominally returned commands. Torque percentages use T_r ; voltage excess is $100 \max(0, V_a/V_M - 1)$. Voltage-violation counts use the reporting allowance stated in the text.

Offset	Cases	Issued	Common	Max. $e_{d,\text{mis}}$ [A]	Max. $e_{q,\text{mis}}$ [A]
10%	Individual	9138	8790	0.9074	1.7219
10%	Joint corners	12,184	11,424	1.2162	2.2111
20%	Individual	9138	8418	1.4423	2.5927
20%	Joint corners	12,184	10,746	3.7653	5.2183
Offset	Cases	Max. $ E_T $ [%]	RMS E_T [%]	Voltage violations	Max. V excess [%]
10%	Individual	17.194	8.044	114 (1.248%)	9.203
10%	Joint corners	32.636	13.933	298 (2.446%)	13.824
20%	Individual	34.388	16.089	217 (2.375%)	18.795
20%	Joint corners	65.272	27.867	607 (4.982%)	28.665

The mismatch errors in Table 10 are substantially larger than the solver errors in Table 9. The largest joint-case torque error occurs for Model 1 with $(\delta_d, \delta_q, \delta_\lambda) = (+0.2, -0.2, -0.2)$ at $\omega = 215$ rad/s and $T_r/T_a = 0.936667$. Here $L_{q,a} - L_{d,a} = 1.152$ mH instead of the nominal 38.93 mH, greatly reducing reluctance torque. The issued current pair is approximately $(-7.06466, 6.83026)$ A; its perturbed-model voltage is 174.33 V, below V_M , but the commanded torque itself is infeasible for that perturbed motor. The 65.27% torque deficit is

therefore an extreme sensitivity result, not a typical saturation estimate or a failure of the scalar optimizer to solve its nominal problem.

Because the frozen controller receives identical nominal inputs in every offset scenario, its boundary-fallback flags are unchanged by construction. Over the same issued-command population, their rates are 42.113% and 18.566% for the low-cost setting in Models 1 and 2, respectively, and 10.714% and 0% for the conservative setting. These rates describe numerical branch selection; they do not detect unknown parameter mismatch. Practical torque and voltage accuracy requires appropriate parameter identification, adaptation, feedback, or independently justified engineering margins. The present test evaluates steady-state model mismatch and does not include current-loop or inverter dynamics.

8.9. Measured Evaluation Burden and Boundary Diagnostics

The added instrumentation counts every scalar residual (F), first derivative (D), and second derivative (H) evaluation during interval construction and objective search, including rejected candidates, acceptance tests, fallback, and inward projection. Coefficient setup, comparisons, and diagnostic quantities evaluated after a solution are excluded. Thus, the count is reproducible but is not processor timing: the three kernels have different arithmetic costs, and shared expressions can be reused in an embedded implementation. The joint-anchor check described in Section 5.2 is included. An audit of all 111,476 setting-specific interval constructions in the parameter-offset study found zero rejections by this additional check, so the parameter-offset results in Section 8.8 are unchanged.

Table 11 uses all 3844 requests of the systematic mixed-objective matrix, including the 773 reference-infeasible requests. The 16-cell and 32-cell direct grids use, respectively, 5 and 6 boundary bisections, and 8 and 12 objective bisections. Both evaluate the same residuals and objective as the proposed method. The errors are $e_d = |\hat{i}_d - i_d^*|$ and $e_q = |\hat{i}_q - i_q^*|$; e_J follows the percentage definition in Section 8.3. Errors are computed on commands for which both the method and the high-accuracy reference return a solution; missed reference-feasible commands are reported separately. None of these methods returned a solution for a reference-infeasible command in this matrix. These baselines are numerical procedures, not implementations of complete previously published drives.

Table 11. Instrumented comparison on 3844 requests (3071 reference feasible). Mean and maximum counts include all requests; errors use the common returned feasible cases. e_J is the relative objective error in percent.

Method	Mean $F+D+H$	Maximum	Max e_d (A)	Max e_q (A)	Max e_J (%)	Misses	Fallback (%)
Low-cost	87.00	178	1.3313×10^{-3}	6.7732×10^{-4}	0.003708	0	29.144
Conservative	91.82	168	2.8920×10^{-4}	1.3305×10^{-4}	0.000341	0	4.852
Grid, 16 cells	58.23	72	1.9474×10^{-2}	1.1368×10^{-2}	0.045298	6	—
Grid, 32 cells	95.76	112	4.5265×10^{-3}	2.3317×10^{-3}	0.006548	1	—

Relative to the 32-cell grid, the mean count decreases by 9.15% for the low-cost setting and 4.12% for the conservative setting. The reported maximum d -axis errors decrease by factors of 3.40 and 15.65, respectively. The proposed method also avoids the narrow-interval misses observed with these grids. The 16-cell grid is cheaper, but has appreciably larger errors and six misses. Both grids have smaller observed maximum counts than the proposed method. This is an accuracy/coverage/cost trade-off, not a claim of universal speed superiority.

Conditioning on the 3071 feasible requests exposes the fallback cost directly. For the low-cost setting, the 2176 requests without fallback average 70.85 evaluations; the 895 fall-

back requests average 119.09. For the conservative setting, the corresponding values are 89.64 (2922 requests) and 142.03 (149 requests). With p the fallback fraction,

$$\mathbb{E}[N \mid \text{feasible}] = (1 - p)\mathbb{E}[N \mid \text{no fallback}] + p\mathbb{E}[N \mid \text{fallback}],$$

giving 84.91 and 92.18 evaluations. Thus, the 29.144% fallback rate retains a mean advantage in this test, but reduces the saving suggested by the fast path alone. The observed worst case is actually larger for the low-cost setting (178 versus 168). Across all requests, acceptance, fallback, and projection contribute, respectively, 7.52, 19.39, and 5.74 mean evaluations for the low-cost setting and 8.12, 6.60, and 5.83 for the conservative setting. Their effort is therefore material and is included in Table 11.

Fallback is not confined to high speed. Pooling the two resistance models, all 216 low-cost Model 1 feasible requests with $T_r/T_a > 0.7$ and $\omega \leq \omega_a$ activate fallback; all 49 such requests with $\omega_a < \omega \leq 2\omega_a$ do so as well. For Model 2 these fractions are 77.78% of 306 and 66.36% of 107 requests, respectively. For Model 1 at intermediate torque ($0.3 < T_r/T_a \leq 0.7$), the fraction is 55.98% of 259 requests between ω_a and $2\omega_a$. Current- and voltage-boundary rejections were distinguished in the numerical analysis. These zones reflect the quality of the finite Taylor/Newton estimates and the strict residual tests; fallback activation by itself is not an indication of current-loop instability or parameter-mismatch detection.

A branch-level analysis also explains Figure 4. The Model 1 maxima occur at $\omega = 278.33$ rad/s for $R_s = 0$ and 262.50 rad/s for full R_s , in both cases at $T_r/T_a = 0.74667$. The voltage candidate fails its root-residual acceptance test and activates voltage-boundary fallback. The constrained optimum is the upper voltage endpoint, so its location error directly becomes the current-reference error: 2.8920×10^{-4} A and 2.2183×10^{-4} A, respectively. The objective solver returns the endpoint without Newton iterations at these peaks. Model 2 has no boundary fallback in the conservative systematic matrix; its maximum error, 2.7234×10^{-5} A, occurs at an interior objective solution and is consistent with its remaining fixed bisection bracket. This diagnosis concerns the numerical procedure for the steady-state model.

From Table 7, one active residual uses at most $2N_B$ residual and $2N_B$ derivative evaluations on the fast path, followed by at most five residual and two derivative acceptance evaluations. If fallback is needed, it adds at most $2N_R + 7$ residual and $N_D + 2$ derivative evaluations. Allowing both constraints, the voltage-inactive checks, and the final joint projection gives conservative caps

$$B_{\text{int}} = 8N_B + 2N_D + 8N_R + 80, \quad B_{\text{obj}} = 2 + 3N_N + N_G.$$

These are upper bounds, not tight simultaneous-path counts. They give interval caps of 248 and 284, and full fixed-command caps of 271 and 314, for the low-cost and conservative settings. The separate full-command caps are

$$N_F \leq 4N_B + 8N_R + 56, \quad N_D^{\text{eval}} \leq 4N_B + 2N_D + 26 + 2N_N + N_G, \quad N_H \leq N_N.$$

Here N_D^{eval} counts derivative evaluations and is distinct from the derivative-bisection iteration budget N_D . The resulting triples are (176, 92, 3) and (200, 110, 4).

For a generic residual $F(x) = Ax^2 + Bx + E + Q/S(x)^2$, precomputed coefficients and a reciprocal-based evaluation require at most six multiplications, four additions, and one reciprocal for F ; five multiplications, three additions, and one reciprocal for F' ; and four multiplications, two additions, and one reciprocal for F'' . Counting these separately bounds the numerical kernels by 1528 multiplications, 986 additions, and 271 reciprocals for the low-cost setting, or 1766, 1138, and 314 for the conservative setting. These arithmetic counts

exclude coefficient construction, comparisons, coordinate transformations, and final i_q recovery. Newton updates, acceptance ratios, and inward correction add at most $4N_B + N_N + 12$ divisions; the two quadratic candidates require at most two square roots. A fused F, F', F'' evaluation can reuse S^{-1} and its powers, reducing the combined kernel to at most 11 multiplications, seven additions, and one reciprocal. Actual machine instructions depend on compiler, division latency, fused operations, and precision; a target-processor cycle measurement is still necessary. Scalar working storage is constant in the iteration budgets.

Branches change the actual execution time, while the loop caps bound the work. A high fallback ratio therefore affects the mean and execution-time distribution without creating an unbounded loop. The optional online torque limiter is more expensive: with one initial request, a zero-torque precheck, N_T interval queries, and one final derated reconstruction, a conservative count is $(N_T + 3)B_{\text{int}} + B_{\text{obj}}$. For the offline verification choice $N_T = 55$, these caps are 14,407 and 16,502 evaluations. Such a choice is not justified for execution at every high-rate PWM interrupt. A slower supervisory update or a smaller torque-resolution budget is appropriate for implementation assessment.

An offline torque map replaces the repeated search by interpolation. A uniform 1D speed table has constant-time cell selection; including measured dc-link voltage requires a two-dimensional table with storage proportional to the number of speed–voltage nodes. Interpolation does not automatically give a conservative feasible torque, and a map generated for nominal voltage or parameters can become stale. A margin and a final feasibility check are still needed. This explains the trade-off between map memory, online adaptation, and the bounded but larger cost of torque bisection.

8.10. Additional Objective Verification and Physical Trade-Off

MTPA and voltage minimization were each tested on the same 31×31 grids for both motors with full R_s : 1922 requested points, of which 1523 are reference feasible. The objectives are the normalized MTPA and voltage indices defined in (47) and (50), respectively, evaluated after torque elimination. Each uses its registered derivatives and initial estimate, with the common interval procedure and unchanged low-cost/conservative budgets. The high-accuracy comparison solves the same corresponding objective and steady-state model. Table 12 reports zero feasibility disagreements for both objectives. This adds numerical evidence of reuse beyond the mixed-index example; it does not validate an arbitrary nonconvex registration or closed-loop drive performance.

Table 12. Additional objective tests, pooled over the two motors with full R_s ; errors and fallback rates use 1523 feasible cases per row.

Objective	Setting	Max e_d (A)	Max e_q (A)	Max e_J (%)	Fallback (%)
MTPA	Low-cost	1.4816×10^{-3}	7.3410×10^{-4}	0.029343	28.956
MTPA	Conservative	3.7506×10^{-4}	2.2420×10^{-4}	0.004422	4.728
Voltage	Low-cost	9.2289×10^{-4}	5.3860×10^{-4}	0.011686	28.956
Voltage	Conservative	2.4703×10^{-4}	1.1948×10^{-4}	0.000721	4.728

Table 13 shows the physical effect at two feasible requests. At the Model 1 point, moving from MTPA to the mixed index increases current magnitude from 4.2764 A to 4.8708 A while reducing voltage demand from 135.4126 V to 113.6171 V. Pure voltage minimization further reduces voltage to 112.2545 V at 5.2542 A. The Model 2 point exhibits the same trade-off. These currents satisfy the same torque equality; the mixed index selects an intermediate current/voltage utilization compromise, without representing a measured loss optimum.

Table 13. Representative objective trade-off with full R_s and the conservative setting. Model 1: $\omega = 300$ rad/s, $T_r = 3.07161$ N m; Model 2: $\omega = 500$ rad/s, $T_r = 2.99136$ N m.

Motor	Objective	i_d (A)	i_q (A)	$ i $ (A)	$ v $ (V)
1	MTPA	−2.0286	3.7646	4.2764	135.4126
1	Mixed	−3.8475	2.9869	4.8708	113.6171
1	Voltage	−4.4485	2.7961	5.2542	112.2545
2	MTPA	−1.9814	3.6099	4.1179	189.2030
2	Mixed	−3.8552	2.8094	4.7703	146.7994
2	Voltage	−5.0855	2.4524	5.6460	139.0604

8.11. DC-Link Voltage Levels and Effective Torque

For an amplitude-invariant dq convention with phase-peak voltages and ideal linear SVPWM, $V_M = V_{dc}/\sqrt{3}$. An implementation may reserve headroom using $V_M = \eta_V V_{dc}/\sqrt{3}$ with $0 < \eta_V < 1$. The added steady-state tests use $\eta_V = 1$ and voltage-limit ratios $r_V \in \{0.8, 1.0, 1.2\}$, giving $V_M = 157.36, 196.70$, and 236.04 V, or equivalent dc-link levels 272.56, 340.69, and 408.83 V. Motor parameters, I_M , the requested torque grid, and nominal $\omega_b = \omega_a$ are held fixed. The mixed index uses the actual V_M in its voltage normalization at each level.

For fixed parameters, speed, and torque, increasing V_M subtracts a larger constant from the voltage residual. Consequently, the exact feasible interval can only expand, and the maximum feasible positive torque cannot decrease. This set-inclusion statement is independent of the selected objective. The optimal current need not vary monotonically because the constraint endpoints and the normalized objective both change. Table 14 reports the feasible count on each 961-request full- R_s grid and representative returned outputs. The conservative setting has no disagreements with the high-accuracy reference on these grids.

Table 14. DC-link-level test with full R_s . The feasible count concerns the full 961-request grid. The remaining columns concern Model 1 at $\omega = 400$ rad/s, $T_r = 6.14322$ N m, or Model 2 at $\omega = 500$ rad/s, $T_r = 4.48704$ N m. T_{max} is the offline reference value; “derated” indicates that the original request was not feasible.

Motor	r_V	Feasible/961	Request Status	T_{max} (N m)	T_r^{eff} (N m)	i_d (A)	i_q (A)
1	0.8	578	Derated	3.34812	3.34477	−4.51523	3.02329
1	1.0	672	Derated	4.46668	4.46221	−5.30879	3.72157
1	1.2	752	Derated	5.70571	5.70000	−6.13699	4.39903
2	0.8	771	Derated	3.47992	3.47644	−5.23726	2.80608
2	1.0	851	Derated	4.35865	4.35429	−4.71365	3.71235
2	1.2	896	Feasible	4.88973	4.48704	−4.57309	3.88418

For an infeasible request, 55 online torque bisections over $[0, T_a]$ use the finite-budget interval procedure; the returned lower feasible estimate is multiplied by $\eta_T = 0.999$ and checked again. The offline T_{max} column instead uses the high-accuracy reference classifier, so the effective torque is not defined by simply multiplying the rounded tabulated maximum by 0.999. At the Model 2 example, the highest voltage level makes the original command feasible and eliminates derating. At the Model 1 example all three commands require derating, but achievable torque increases with voltage. These are independent steady-state operating points, not a transient dc-link ripple simulation. A measured voltage drop in a drive therefore requires an updated feasibility computation and appropriate dynamic voltage reserve, rather than continued use of references or torque maps calculated at nominal voltage.

8.12. Tolerance Sensitivity and Near-Tangent Classification

We tested three midpoint/root multiplier pairs in (41): $(10^{-12}, 10^{-14})$, $(10^{-10}, 10^{-12})$, and $(10^{-8}, 10^{-10})$. For each pair, $\varepsilon_B/I_M \in \{10^{-4}, 10^{-3}, 10^{-2}\}$ was tested with both solver settings on all 3844 systematic requests, for 69,192 runs. The iteration budgets, coefficient values, derivative zero guards, and high-accuracy reference were unchanged. Table 15 reports the results. For each row the three boundary-location tolerances give the same reported maximum errors, classifications, and fallback percentage; this agreement is empirical for this matrix, not a general insensitivity theorem.

Table 15. Sensitivity to paired midpoint/root relative multipliers in double precision. All rows have zero misses and zero false-feasible decisions on the systematic grid. Each row represents separate tests at $\varepsilon_B/I_M = 10^{-4}, 10^{-3}, 10^{-2}$ with identical reported aggregates.

Multipliers	Setting	Max e_d (A)	Max e_q (A)	Fallback (%)
$(10^{-12}, 10^{-14})$	Low-cost	1.3313×10^{-3}	6.7732×10^{-4}	34.256
$(10^{-10}, 10^{-12})$	Low-cost	1.3313×10^{-3}	6.7732×10^{-4}	29.144
$(10^{-8}, 10^{-10})$	Low-cost	1.3313×10^{-3}	6.7732×10^{-4}	25.887
$(10^{-12}, 10^{-14})$	Conservative	2.8920×10^{-4}	1.3305×10^{-4}	7.392
$(10^{-10}, 10^{-12})$	Conservative	2.8920×10^{-4}	1.3305×10^{-4}	4.852
$(10^{-8}, 10^{-10})$	Conservative	2.8920×10^{-4}	1.3305×10^{-4}	3.680

Tighter residual tests increase fallback even though the final grid errors do not change here; loosening them reduces this work. The root-residual condition is more restrictive than the tested ε_B choices in the cases governing these aggregates. The default $10^{-3}I_M$ is therefore a reproducible candidate-screening choice used together with the stricter relative residual test and final feasibility correction, not a uniquely optimal tolerance. Scaling and floating-point precision must be considered jointly before transferring it.

To avoid interpreting a regular grid as evidence of exact boundary classification, additional full- R_s probes use $T_r = T_{\max}^{\text{ref}}(1 + \delta)$ with $\delta \in \{-10^{-4}, -10^{-6}, -10^{-8}, 0, 10^{-8}, 10^{-6}, 10^{-4}\}$, at 300, 400, and 500 rad/s for Model 1 and 500, 600, and 700 rad/s for Model 2. Each of the 18 tolerance/setting combinations therefore has 42 probes, including 24 reference-feasible probes. The low-cost setting returns 14 of these and the conservative setting returns 20; neither returns a reference-infeasible point. These aggregate counts are unchanged across the tested acceptance constants. The conservative setting resolves all 18 strictly submaximum probes, but only two of the six numerical $T_{\max}^{\text{ref}}$ points. The low-cost setting also leaves four strictly submaximum probes unresolved. Near a double boundary root, a finite derivative-minimum search can fail to locate the very narrow negative residual region, and a tolerance-accepted interval may lack a joint feasible midpoint. The additional anchor guard explicitly declines the latter case before attempting projection. Thus, the finite-budget predicate is conservative in these tests but is not identical to exact feasibility arbitrarily close to tangency. Further refinement or a derated request is needed when that boundary resolution matters.

8.13. Drive Integration, Dynamic Transfer, and Signed Operation

Figure 5 places the generator between a torque/speed supervisor and the inner current loops. Its inputs are the requested torque, measured/estimated electrical speed, dc-link-derived voltage budget, current limit, and adopted motor parameters. It returns the current pair only after the implemented checks; effective torque and status are also returned to the supervisor for saturation handling. The diagram specifies the interface required for a real driver without representing an experimentally validated implementation of the complete system.




Figure 5. Proposed integration into a field-oriented motor driver. The thick generator block is the contribution studied here; the surrounding regulator, modulation, sensing, and supervisory functions are required for dynamic operation. Speed and dc-link voltage are refreshed from the drive measurements.

For constant inductances, the dynamic voltage additionally includes

$$v_{d,\text{dyn}} = v_{d,\text{ss}} + L_d \dot{i}_d, \quad v_{q,\text{dyn}} = v_{q,\text{ss}} + L_q \dot{i}_q.$$

A reference lying on the steady-state voltage boundary may therefore leave insufficient voltage for a fast current transient. A practical integration needs voltage reserve, reference slew limits, current-regulator anti-windup, and coordination with torque saturation. Filtering or interpolating successive current pairs does not automatically preserve their original torque equality or feasibility at the updated speed and voltage; the conditioned request must be checked again. Such an algebraic check still does not prove dynamic reachability or closed-loop stability.

The reference layer can run at a slower supervisory rate than the inner current loop if speed/voltage changes between updates are covered by suitable reserves. Sampling and computation delay, PWM saturation, dead time, measurement noise, parameter drift, and reduced current-loop bandwidth at high speed all affect performance. An unresolved status must be handled by the supervisor with a checked derated request or the drive's defined protective state. Regulator gains, update rates, and protective actions depend on the actual drive and have not been validated by the present static tests. A low scalar count alone cannot establish a high sampling-rate capability without target-hardware timing and closed-loop verification.

Signed Torque and Speed on the Same d -Axis Branch

Under the same constant-parameter assumptions $L_q > L_d > 0$, $\lambda_m > 0$, and $i_d \leq 0$, retain signed $C = T_r/K_t$ and signed ω in $i_q = C/S(i_d)$. Then i_q has the sign of torque. In the full-resistance voltage residual, $A = R_s^2 + \omega^2 L_d^2 \geq 0$ and $Q = (R_s^2 + \omega^2 L_q^2)C^2 \geq 0$ remain nonnegative for either sign, while the constant term contains the signed cross term $2R_s\omega C$. The same second-derivative expressions for the current and voltage residuals therefore remain nonnegative. Each fixed signed torque–speed problem on this i_d branch still has a connected feasible interval. Jointly reversing (C, ω) preserves the d -axis interval and reverses i_q ; changing only one sign must retain the changed resistive cross term.

This extension of interval geometry does not by itself extend the original torque-magnitude nesting proof to regeneration. Writing $C = \sigma q$ with $q \geq 0$ and $\sigma \in \{-1, 1\}$ gives

$$\frac{\partial F_{V,R}}{\partial q} = 2R_s\omega\sigma + \frac{2(R_s^2 + \omega^2 L_q^2)q}{S(i_d)^2}.$$

For motoring, $\omega\sigma \geq 0$, whereas regeneration makes the first term negative; hence the original pointwise nonnegative-derivative argument cannot be reused without further analysis. For $R_s = 0$ that obstruction disappears. A four-quadrant implementation needs signed current bounds and outputs, quadrant-aware torque limiting with an established search property, continuity/hysteresis through zero torque or speed, and a dc-link energy-management strategy during regeneration. Simply applying absolute values to torque and speed is insufficient for the full- R_s model. The present numerical validation remains restricted to positive motoring; no four-quadrant closed-loop result is inferred from this algebraic extension. Nonlinear magnetic maps can invalidate these convexity arguments, as discussed in Section 8.14.

8.14. Limitations and Future Validation

The principal limitations are as follows.

First, L_d , L_q , R_s , and λ_m are constant throughout each scalar optimization. The offset tests in Section 8.8 preserve this assumption and quantify coefficient sensitivity and frozen-model mismatch; they do not reproduce current-dependent saturation or cross-coupling. A constant parameter update satisfying Assumption 1 preserves the stated convexity and connected-interval results for the updated model. It does not, by itself, certify the physical machine when the supplied parameters are inaccurate. With nonlinear flux maps $\psi_d(i_d, i_q)$ and $\psi_q(i_d, i_q)$, magnetic saturation and cross-coupling can change the torque and voltage geometry [25]. In that setting, the explicit elimination (9) and the constant-coefficient second-derivative formulas need not remain valid. Even when a local torque-equality branch can be solved for i_q as a function of i_d , the resulting current and voltage residuals need not be convex, and the feasible set need not be one connected interval. Proposition 2's monotone positive-torque feasibility also requires a new justification. Evaluating a flux map and freezing its parameters during each call retains the guarantees only for that local surrogate; convergence and physical feasibility of repeated map updates require separate analysis. Direct nonlinear-map use therefore requires verification of branch coverage, feasibility geometry, and objective structure, or an alternative constrained search that covers every admissible component. The present theorem is not extended to arbitrary nonlinear magnetic characteristics.

Second, the main theorem and numerical validation use the positive-motoring branch. Section Signed Torque and Speed on the Same d -Axis Branch gives the constant-parameter signed extension of the interval geometry, but regenerative torque limiting, sign transitions, dc-link energy handling, and four-quadrant closed-loop operation still require separate implementation and validation.

Third, the study is steady-state. The added dc-link cases are independent voltage levels; they do not simulate dc-link ripple, current-loop dynamics, dead time, sampling delay, PWM saturation, or reference smoothing. The block diagram identifies the required surrounding drive functions without claiming that their closed-loop performance has been demonstrated.

Fourth, the proposed current–flux and parameter-error objectives are design examples, not experimentally calibrated thermal or reliability metrics. Fifth, no processor timing, memory measurement, or hardware experiment is reported.

Future work should therefore integrate saturation-aware maps, inverter nonlinearity compensation, parameter adaptation, signed four-quadrant branch handling, and transition smoothing. Hardware tests should measure constraint adherence, current-reference tracking, computational deadline margin, and loss/temperature outcomes when a physically identified loss or thermal index is used.

9. Conclusions

This study reorganized positive-motoring IPMSM current-reference generation into exact torque elimination, a connected current/voltage-feasible interval, and safeguarded scalar performance-index minimization. The formulation removes the zero-speed singularity from the simplified voltage residual, proves monotone torque feasibility, includes a zero-torque precheck, and prevents unchecked Taylor/Newton roots from defining the interval. For objectives admitted to the certified strictly convex path, a derivative-sign bracket and mandatory bisection contractions provide an explicit final-bracket error bound, while optional Newton proposals accelerate contraction. Across 3071 feasible points in a 3844-point systematic sweep, the low-cost and conservative settings produced maximum d -axis current errors of 0.001331 and 0.000289 A, respectively, relative to a high-accuracy scalar reference. MTPA, fixed-torque voltage minimization interpreted as the dual formulation of MTPV, loss-oriented objectives, a mixed current-voltage index, and two structurally convex extensions use the same feasibility layer; conditional or lookup-table objectives are routed to a registered bounded-domain scalar fallback, without a global-optimality claim unless an appropriate certificate is available. Additional 10–20% constant-parameter tests retained small numerical solution errors when the model coefficients were updated, while a frozen nominal controller exhibited appreciable steady-state torque and voltage errors. The numerical safeguards therefore provide reliability of the solution procedure under the adopted model; they are not a substitute for parameter adaptation or drive-level robustness validation. On the 3844-command comparison, the low-cost and conservative settings required 87.00 and 91.82 mean scalar residual/derivative evaluations, versus 95.76 for the 32-cell grid baseline. Their maximum d -axis errors were 3.40 and 15.65 times smaller, respectively, and neither missed a reference-feasible command, whereas the 32-cell grid missed one. These gains include the measured acceptance, fallback, and projection work, but come with larger observed worst-case counts of 178 and 168 versus 112. Thus, the quantitative advantage is the tested accuracy and narrow-interval detection at a competitive mean evaluation burden, with explicit bounded recovery; it is not universal superiority in runtime or worst-case cost. Additional MTPA and voltage-minimization tests support reuse of the same feasibility layer. The voltage-level and near-tangent tests also make the attainable torque and finite-budget classification limits explicit. These results support the framework as a reproducible, fixed-budget numerical current-reference layer under the stated assumptions. Hardware timing, nonlinear magnetic maps, four-quadrant operation, and experimental validation remain necessary before drive-level real-time or efficiency claims are made.

Supplementary Materials: The following supporting information can be downloaded at <https://www.mdpi.com/article/10.3390/electronics15184313/s1>: The Base-MATLAB script, compatible with MATLAB R2024a or later, used to reproduce Figures 1–4 and generate their associated plotting data.

Author Contributions: Conceptualization, H.H.C.; methodology, H.H.C.; software, H.H.C.; validation, D.K. and H.H.C.; formal analysis, H.H.C.; investigation, D.K. and H.H.C.; writing—original draft preparation, H.H.C.; writing—review and editing, D.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 Korean government’s Ministry of Trade, Industry and Energy under the project titled, “Middle market enterprise specialized human resources development for residential and commercial fuel cell”, under number 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 1–4 and generate their associated plotting data is provided as Supplementary Material.

Acknowledgments: During the preparation of this manuscript, the authors used OpenAI ChatGPT 5.6 for English-language editing, LaTeX restructuring, equation cross-checking, and assistance in preparing reproducibility scripts. The authors independently verified all equations, code, numerical results, references, and interpretations; reviewed and edited the output; and take full responsibility for the content of this publication.

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

Abbreviations

The following abbreviations are used in this manuscript:

1D	One-dimensional
FW	Field weakening
IPMSM	Interior permanent-magnet synchronous motor
LMC	Loss-minimization control
MTPA	Maximum torque per ampere
MTPV	Maximum torque per voltage
PMSM	Permanent-magnet synchronous motor

References

1. 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. [\[CrossRef\]](#)
2. 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. [\[CrossRef\]](#)
3. Morimoto, S.; Tong, Y.; Takeda, Y.; Hirasaka, T. Loss minimization control of permanent magnet synchronous motor drives. *IEEE Trans. Ind. Electron.* **1994**, *41*, 511–517. [\[CrossRef\]](#)
4. Krishnan, R. *Permanent Magnet Synchronous and Brushless DC Motor Drives*; CRC Press: Boca Raton, FL, USA, 2010.
5. Sul, S.-K. *Control of Electric Machine Drive Systems*; Wiley-IEEE Press: Hoboken, NJ, USA, 2011.
6. Jung, S.-Y.; Hong, J.; Nam, K. Current minimizing torque control of the IPMSM using Ferrari’s method. *IEEE Trans. Power Electron.* **2013**, *28*, 5603–5617. [\[CrossRef\]](#)
7. 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. [\[CrossRef\]](#)
8. Choi, K.; Kim, Y.; Kim, K.-S.; Kim, S.-K. Real-time optimal torque control of interior permanent magnet synchronous motors based on a numerical optimization technique. *IEEE Trans. Control Syst. Technol.* **2021**, *29*, 1815–1822. [\[CrossRef\]](#)
9. Miguel-Espinar, C.; Heredero-Peris, D.; Villafafila-Robles, R.; Montesinos-Miracle, D. Review of flux-weakening algorithms to extend the speed range in electric vehicle applications with permanent magnet synchronous machines. *IEEE Access* **2023**, *11*, 22961–22981. [\[CrossRef\]](#)
10. 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. [\[CrossRef\]](#)
11. Dianov, A.; Tinazzi, F.; Calligaro, S.; Bolognani, S. Review and classification of MTPA control algorithms for synchronous motors. *IEEE Trans. Power Electron.* **2022**, *37*, 3990–4007. [\[CrossRef\]](#)
12. Huang, K.; Peng, W.; Lai, C.; Feng, G. Efficient maximum torque per ampere (MTPA) control of interior PMSM using sparse Bayesian based offline data-driven model with online magnet temperature compensation. *IEEE Trans. Power Electron.* **2023**, *38*, 5192–5203. [\[CrossRef\]](#)
13. Xu, Y.; Lin, M.-Y.; Wan, W.-J.; Wang, Y.; Shen, J.-X. Maximum torque per ampere control of PMSMs using real-time numerical optimization. *IEEE Access* **2025**, *13*, 118611–118624. [\[CrossRef\]](#)

14. Zhang, H.; Xiao, D. Steady-state efficiency optimization method for PMSM considering current distribution and iron loss resistance. *IEEE Trans. Circuits Syst. II Express Briefs* **2024**, *71*, 4924–4928. [[CrossRef](#)]
15. Liu, Z.; Mao, K.; Lei, X.; Zheng, S.; Zhang, H. Loss minimization control based on bivariate extreme value theory for PMSMs. *IEEE Trans. Power Electron.* **2024**, *39*, 2004–2012. [[CrossRef](#)]
16. Jung, B.; Lee, J.; Choi, J.; Kim, Y. Analytical loss minimization control for IPMSM using linearized torque and loss modeling. *IEEE Access* **2026**, *14*, 2224–2235. [[CrossRef](#)]
17. Liu, L.; Guo, Y.; Yin, W.; Lei, G.; Sun, X.; Zhu, J. Efficient design optimization of PMSM drive systems using improved equivalent-circuit-based loss minimization control. *IEEE Trans. Ind. Electron.* **2025**, *72*, 3280–3291. [[CrossRef](#)]
18. Hou, Q.; Wang, H.; Lee, C.H.T.; Ding, S. Composite adaptive super-twisting sliding mode control using barrier function for PM motor drives toward electric aircraft applications. *IEEE Trans. Power Electron.* **2025**, *40*, 16255–16264. [[CrossRef](#)]
19. Hou, Q.; Wang, H.; Zhao, C.; Xu, S.; Zuo, Y.; Lee, C.H.T.; Ding, S. Super-twisting extended state observer-based quasi-proportional-resonant controller for permanent magnet synchronous motor drive system. *IEEE Trans. Transp. Electr.* **2024**, *10*, 1596–1604. [[CrossRef](#)]
20. Kim, S.-M.; Choi, H.H. Low-complexity residual-corrected loss-minimization current-reference generation for PMSM drives. *Electronics* **2026**, *15*, 3000. [[CrossRef](#)]
21. Kang, D.; Choi, H.H. Low-complexity SVS current-reference generation for PMSMs with and without MTPV endpoints. *Energies* **2026**, *19*, 3621. [[CrossRef](#)]
22. Dai, L.; Niu, S.; Yuan, X. Current-injected torque ripple reduction with minimized iron loss in PMSMs. *IEEE Trans. Ind. Electron.* **2025**, *72*, 9790–9794. [[CrossRef](#)]
23. Sun, X.; Su, Z.; Lei, G.; Yao, M. Robust predictive cascaded speed and current control for PMSM drives considering parameter variations. *IEEE Trans. Ind. Electron.* **2024**, *71*, 10235–10245. [[CrossRef](#)]
24. 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](#)]
25. Stumberger, B.; Stumberger, G.; Dolinar, D.; Hamler, A.; Trlep, M. Evaluation of saturation and cross-magnetization effects in interior permanent-magnet synchronous motor. *IEEE Trans. Ind. Appl.* **2003**, *39*, 1264–1271. [[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.