



RESEARCH ARTICLE

Conformally Calibrated Residual-Learning Guidance for Field-of-View Constrained High-Speed Interception

Trinh Thi Minh¹, Vu Quang Luong^{2*}

¹Basic Sciences Faculty, Air Defence-Air Force Academy, Hanoi, Vietnam

²Missile Faculty, Air Defence-Air Force Academy, Hanoi, Vietnam

¹ORCID: <https://orcid.org/0009-0002-6863-7369>

²ORCID: <https://orcid.org/0009-0003-6599-1159>

*Corresponding Author: vuquangluong88@gmail.com

ABSTRACT

This paper proposes a certifiable residual learning-enhanced guidance law for high-speed interception subject to field-of-view constraints, measurement delay, and actuator dynamics. A residual predictor estimates the target-maneuver component that is not represented by the physics-based model. Its prediction error is calibrated through split conformal prediction and embedded directly into a robust high-order control-barrier-function condition. The final acceleration command is generated by a CLF–HOCBF safety filter that minimally modifies the nominal command when constraint satisfaction is endangered. The predictor is constructed from 87,560 simulated samples and achieves a held-out RMSE of 5.768 m/s² with 94.83% empirical coverage for a nominal 95% marginal target. The resulting conformal radius is 12.194 m/s², which is 51.16% smaller than the fixed comparison bound. Over 220 Monte Carlo trials for each method and parameter set, the proposed method reduces the median and 95th-percentile terminal errors by 32.64% and 35.26%, respectively, under the moderate condition relative to physics-only guidance. Under the severe near-boundary condition, it reduces the 95th-percentile error by 8.60%, operates at a median field-of-view utilization of 96.64%, and produces no observed field-of-view violation within the investigated domain. The results demonstrate improved maneuver compensation together with an auditable, data-calibrated uncertainty margin for safety filtering.

Keywords: conformal prediction, control barrier function, field-of-view constraint, high-speed interception, residual learning, solver-free safety filter

1. INTRODUCTION

Physics-based guidance laws remain attractive because of their transparent structure, analyzable stability properties, and moderate real-time computational requirements. During high-speed terminal interception, however, guidance performance is simultaneously affected by unknown target acceleration, delayed and noisy measurements, inner-loop dynamics, command saturation, and the finite field of view of the seeker. Integrated guidance-and-control methods have addressed terminal-angle and field-of-view constraints, but their performance still depends on the assumed maneuver model and prescribed uncertainty bounds [1].

Machine-learning methods have been introduced to exploit nonlinear relationships between measurement histories and target maneuvers. Reinforcement-learning guidance has been investigated for angle-only interception [2], all-aspect attack [3], field-of-view-constrained impact-time control [4], proximal-policy-optimization guidance [5], and engagements with time-varying velocity [6]. These studies demonstrate that learned models can improve performance within selected training and evaluation domains. Nevertheless, a low test error does not constitute a reliable online uncertainty bound, and a constraint penalty in the training objective does not establish forward invariance when the engagement departs from the training distribution.

Control barrier functions convert state constraints into inequalities imposed on the control input and can be

implemented through quadratic programs [7]. High-order control barrier functions extend this framework to constraints with relative degree greater than one [8]. Adaptive CBF and barrier-based safe-reinforcement-learning methods further address parametric uncertainty and learned policies [9]–[11]. Their validity, however, depends on the uncertainty bound used in the robust constraint. An excessively large bound introduces unnecessary conservatism, whereas an underestimated bound invalidates the claimed certificate.

Conformal prediction constructs finite-sample marginal coverage statements from residuals evaluated on an independent calibration set without requiring a Gaussian error model [12], [13]. This property is attractive for learning-enhanced guidance because the uncertainty radius can be derived from data rather than selected heuristically. Conformal calibration alone, however, does not enforce dynamical constraints. The remaining research gap is a unified architecture that preserves a physics-based guidance structure, learns only the unmodeled maneuver component, calibrates the learning error from data, embeds the calibrated radius into the field-of-view condition, and simultaneously accounts for measurement delay, actuator-tracking error, and command limits in a real-time solution.

To address this gap, the present study develops a conformal-certified residual learning-enhanced guidance law with robust CLF–HOCBF safety filtering. The residual predictor estimates the unmodeled target acceleration, the split-conformal quantile

enters the robust barrier inequality, and the resulting scalar safety-filter problem is solved through closed-form projection and candidate evaluation. The scientific contribution is therefore the integration of data-driven maneuver compensation, finite-sample uncertainty quantification, and dynamical constraint enforcement within a single interpretable and reproducible guidance architecture.

The novelty lies in the way these elements are coupled rather than in claiming any constituent technique as new in isolation. In contrast to end-to-end learned guidance, learning is restricted to the unmodeled maneuver residual while the nominal feedback structure remains physics based. In contrast to standalone conformal regression, the calibrated residual radius is propagated into the dynamical HOCBF inequality. In contrast to robust CBF designs that use a single hand-selected uncertainty margin, the learning-error radius is data calibrated and kept separate from deterministic geometry and actuator envelopes. Finally, because the command is scalar, the CLF-HOCBF filter is reduced to exact interval projection and finite candidate evaluation, avoiding dependence on a generic online quadratic-program solver.

2. METHODOLOGY

2.1. Relative-Motion Model and Constraints

The line-of-sight angle is propagated by its angular rate:

$$\dot{q} = \omega \quad (1)$$

Here, q is the line-of-sight angle and ω is its angular rate.

With positive closing speed, the angular-rate dynamics are written as:

$$\dot{\omega} = \frac{2V_c}{R}\omega + \frac{a_T - a_M + d_g}{R} \quad (2)$$

Here, R is relative range, V_c is closing speed, a_T is target lateral acceleration, a_M is realized interceptor acceleration, and d_g is reduced geometric/model mismatch.

The inner acceleration loop is represented by a saturated first-order model:

$$\tau_a \dot{a}_M = -a_M + \text{sat}_{[u_{\min}, u_{\max}]}(u) + d_a \quad (3)$$

Here, τ_a is the actuator time constant, u is the safety-filtered command, d_a is tracking disturbance, and $u_{\min}$ and $u_{\max}$ are the command limits.

By defining the inner-loop tracking error, the reduced state dynamics can be separated into nominal and uncertain parts:

$$\mathbf{x} = [q \ \omega]^T, \dot{\mathbf{x}} = \begin{bmatrix} \omega \\ \frac{2V_c}{R}\omega - \frac{u}{R} \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{a_T + d_g - e_a}{R} \end{bmatrix}, e_a = a_M - u \quad (4)$$

Here, $\mathbf{x}$ is the reduced relative state and e_a is the inner-loop tracking error.

The seeker constraint and admissible command set are therefore:

$$C = \{\mathbf{x} : h(\mathbf{x}) = q_{\max}^2 - q^2 \geq 0\}, \quad u_{\min} \leq u \leq u_{\max} \quad (5)$$

Here, h is the field-of-view function, $q_{\max}$ is the permitted half-angle, and C is the safe set.

2.2. Physics-Preserving Residual Learning

The nominal command retains the physical guidance structure and uses the learned maneuver estimate only as feedforward compensation:

$$u_L = \hat{a}_T + 2V_c\omega + R(k_p q + k_d \omega) \quad (6)$$

Here, u_L is the unfiltered nominal command, k_p and k_d are positive feedback gains, and $\hat{a}_T$ is the residual predictor output.

Delayed acceleration measurements are low-pass filtered and assembled into a normalized short-history feature vector:

$$\bar{a}_k = \gamma \bar{a}_{k-1} + (1-\gamma)a_{k-d}^m, \quad \xi_k = \mathbf{D}_\xi^{-1} \begin{bmatrix} \bar{a}_k & \bar{a}_{k-1} & \bar{a}_{k-2} & \frac{\bar{a}_k - \bar{a}_{k-1}}{T_s} & \frac{\bar{a}_k - 2\bar{a}_{k-1} + \bar{a}_{k-2}}{T_s^2} & r_k & \frac{\tau_{go,k}}{\tau_{go,0}} \end{bmatrix}^T \quad (7)$$

Here, $\bar{a}$ is the filtered measurement, γ is its smoothing factor, a_m is the measured proxy, d is the delay in samples, T_s is the sampling period, ξ_k is the feature vector, $\mathbf{D}_\xi$ is a diagonal normalizer, r_k is normalized range, and τ_{go} is time to go.

The compact random-feature model and its scalar maneuver prediction are:

$$\phi_k = \begin{bmatrix} \tanh(\mathbf{W}_h \xi_k + \mathbf{b}_h) \\ 1 \end{bmatrix}, \quad \hat{a}_{T,k} = s_y \phi_k^T \mathbf{w}_o \quad (8)$$

Here, ϕ_k is the augmented hidden feature, $\mathbf{W}_h$ and $\mathbf{b}_h$ are fixed hidden parameters, s_y is the output scale, and $\mathbf{w}_o$ is the trained output-weight vector.

The output layer is obtained in one ridge-regression step:

$$\mathbf{w}_o = (\Phi^T \Phi + \lambda_r \mathbf{I})^{-1} \Phi^T \mathbf{y} \quad (9)$$

Here, Φ stacks the hidden features, $\mathbf{y}$ stacks the training targets, λ_r is the ridge coefficient, and $\mathbf{I}$ is the identity matrix.

2.3. Split-Conformal Error Calibration

An independent calibration subset produces absolute nonconformity scores:

$$s_i = |a_{T,i} - \hat{a}_{T,i}|, \quad i = 1, \dots, n_{cal} \quad (10)$$

Here, s_i is the i^{th} calibration score and n_{cal} is the number of calibration samples.

For a declared miscoverage level, the finite-sample order statistic is selected as:

$$k_\alpha = \lceil (n_{cal} + 1)(1 - \alpha) \rceil, \quad \bar{e}_\alpha = s_{(k_\alpha)} \quad (11)$$

Here, α is the target miscoverage, k_α is the order-statistic index, $s_{(k_\alpha)}$ is the ordered score, and $\bar{e}_\alpha$ is the calibrated residual radius.



Under exchangeability between calibration and future residuals, the resulting interval satisfies:

$$\Pr\{|a_r - \hat{a}_r| \leq \bar{e}_\alpha\} \geq 1 - \alpha \quad (12)$$

This is a marginal finite-sample statement; it is not a conditional guarantee at every state.

The implementation randomly partitions feature-target rows after the maneuver trajectories have been generated. Residuals originating from the same trajectory can therefore be temporally correlated across the partitions, so exact exchangeability of closed-loop time samples is not established by this experiment. Equation (12) is consequently interpreted as a marginal calibration statement for new samples drawn from the declared simulation generator, not as a statewise guarantee or as a distribution-free guarantee under arbitrary deployment shift. If the maneuver process, sensor statistics, measurement delay, or actuator dynamics depart from the calibration regime, the radius must be re-estimated using trajectory-level or block splits, weighted conformal methods, or online recalibration before the safety claim is invoked. The HOCBF certificate remains conditional on the conformal event and on the deterministic error envelopes stated below.

2.4. Robust High-Order Field-of-View Barrier

Because the field-of-view constraint has relative degree two with respect to the acceleration command, the first two barrier variables are:

$$\begin{aligned} \psi_0 &= h = q_{\max^2} - q^2 \\ \psi_1 &= \dot{h} + \alpha_1 \psi_0 = -2q\omega + \alpha_1 (q_{\max^2} - q^2) \end{aligned} \quad (13)$$

Here, ψ_0 and ψ_1 are the zeroth- and first-order barrier variables, and α_1 is a positive barrier gain.

Differentiating ψ_1 and substituting the relative-motion model gives the complete second-order expression:

$$\begin{aligned} \psi_2 &= -2\omega^2 - \frac{4V_c}{R} q\omega - \frac{2q}{R} (\hat{a}_r - u + \tilde{e}_r + d_g - e_a) \\ &\quad - 2\alpha_1 q\omega + \alpha_2 [-2q\omega + \alpha_1 (q_{\max^2} - q^2)] \end{aligned} \quad (14)$$

Here, ψ_2 is the second-order barrier variable, $\tilde{e}_r$ is the prediction residual, and α_2 is the second positive barrier gain.

The command-dependent and uncertain terms are separated as:

$$\psi_2 = \frac{2q}{R} u + \beta_H(\mathbf{x}, R, V_c, \hat{a}_r) - \frac{2q}{R} (\tilde{e}_r + d_g - e_a) \quad (15)$$

Here, β_H collects all known state, geometry, and predicted-maneuver terms.

For direct implementation, β_H is expanded without hidden terms:

$$\begin{aligned} \beta_H &= -2\omega^2 - \frac{4V_c}{R} q\omega - \frac{2q}{R} \hat{a}_r - 2\alpha_1 q\omega \\ &\quad + \alpha_2 [-2q\omega + \alpha_1 (q_{\max^2} - q^2)] \end{aligned} \quad (16)$$

No new symbol is introduced in this expansion.

Bounding the prediction, geometry, and actuator errors yields

the sufficient robust HOCBF condition:

$$\frac{2q}{R} u + \beta_H - \frac{2|q|}{R} (\bar{e}_\alpha + \bar{d}_g + \bar{e}_a) \geq 0 \quad (17)$$

Here, $\bar{d}_g$ and $\bar{e}_a$ are deterministic bounds on geometric mismatch and inner-loop error. The conformal radius $\bar{e}_\alpha$ is the only data-calibrated term.

For the reported normalized simulations, the propagation model and controller use the same reduced relative-motion geometry; hence the declared geometric/model-mismatch envelope is set to zero within this specific simulation domain. The actuator-tracking allowance is selected a priori as 0.12 times the command limit, corresponding to $\text{CFG.innerAllowanceFraction} = 0.12$, and is added to the conformal residual radius before normalization by range. This allowance is not fitted to the Monte Carlo outcomes. For transfer to a physical system, the geometric envelope must instead be estimated from independent model-validation residuals, and the actuator envelope must be established from closed-loop tracking tests over the admissible command and command-rate range. The certificate applies only while the realized errors remain within these declared envelopes.

2.5. Robust Control-Lyapunov Constraint

The nominal closed-loop matrix and its Lyapunov matrix are defined by:

$$\mathbf{A}_c = \begin{bmatrix} 0 & 1 \\ -k_p & -k_d \end{bmatrix}, \quad \mathbf{A}_c^T \mathbf{P} + \mathbf{P} \mathbf{A}_c = -\mathbf{I}, \quad \mathbf{P} = \mathbf{P}^T \succ 0 \quad (18)$$

Here, $\mathbf{A}_c$ is the nominal error-dynamics matrix and $\mathbf{P}$ is the symmetric positive-definite solution of the Lyapunov equation.

The convergence measure is chosen as the quadratic function:

$$V(\mathbf{x}) = \mathbf{x}^T \mathbf{P} \mathbf{x} \quad (19)$$

Here, V is the control-Lyapunov function.

Its derivative is expanded into nominal and uncertainty contributions:

$$\dot{V} = 2\mathbf{x}^T \mathbf{P} \begin{bmatrix} \omega \\ \frac{2V_c}{R} \omega + \frac{\hat{a}_r - u}{R} \end{bmatrix} + \frac{2\mathbf{x}^T \mathbf{P} \mathbf{e}_2}{R} (\tilde{e}_r + d_g - e_a), \quad \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (20)$$

Here, $\mathbf{e}_2$ selects the angular-rate channel.

Using the same uncertainty bounds gives the softened robust CLF inequality:

$$L_f V + L_g V u + \frac{2|\mathbf{e}_2^T \mathbf{P} \mathbf{x}|}{R} (\bar{e}_\alpha + \bar{d}_g + \bar{e}_a) \leq -c_v V + \delta \quad (21)$$

Here, $L_f V$ and $L_g V$ are the Lie-derivative terms, c_v is the requested decay rate, and δ is a nonnegative relaxation variable.

2.6. Solver-Free Lyapunov-Barrier Projection

At each sample, the final command minimizes the deviation from the physics-learning command:



$$\min_{u, \delta \geq 0} \frac{1}{2}(u - u_L)^2 + \frac{p_\delta}{2} \delta^2 \quad (22)$$

s.t. robust HOCBF, robust CLF, and $u_{\min} \leq u \leq u_{\max}$

Here, p_δ is the penalty applied to the CLF relaxation. The HOCBF and actuator bounds remain hard constraints.

The robust HOCBF inequality is first written in scalar affine form:

$$\begin{aligned} a_H u &\geq b_H, \quad a_H = \frac{2q}{R} \\ b_H &= -\beta_H + \frac{2|q|}{R}(\bar{e}_\alpha + \bar{d}_g + \bar{e}_a) \end{aligned} \quad (23)$$

Here, a_H and b_H are the barrier coefficient and threshold.

The hard field-of-view and command constraints define the feasible scalar interval:

$$\begin{aligned} \ell_H &= \begin{cases} b_H / a_H, & a_H > 0 \\ -\infty, & a_H \leq 0 \end{cases} \\ r_H &= \begin{cases} b_H / a_H, & a_H < 0 \\ +\infty, & a_H \geq 0 \end{cases} \\ \underline{u} &= \max(u_{\min}, \ell_H), \quad \bar{u} = \min(u_{\max}, r_H) \end{aligned} \quad (24)$$

Here, ℓ_H and r_H are the sign-dependent barrier bounds, while $\underline{u}$ and $\bar{u}$ are the final interval limits.

After eliminating the optimal CLF slack, the scalar objective becomes a hinge-penalized quadratic:

$$J(u) = \frac{1}{2}(u - u_L)^2 + \frac{p_\delta}{2} [\max(0, a_C u - b_C)]^2, \quad u \in [\underline{u}, \bar{u}] \quad (25)$$

Here, a_C and b_C are the affine CLF coefficient and threshold, and J is the reduced one-dimensional cost.

The stationary point and finite candidate set provide the exact solver-free command:

$$\begin{aligned} u_C &= \frac{u_L + p_\delta a_C b_C}{1 + p_\delta a_C^2}, \\ S &= \{\underline{u}, \bar{u}, \Pi_{[\underline{u}, \bar{u}]}(u_L), \Pi_{[\underline{u}, \bar{u}]}(b_C / a_C), \Pi_{[\underline{u}, \bar{u}]}(u_C)\}, \\ u^\ddagger &= \arg \min_{u \in S} J(u) \end{aligned} \quad (26)$$

Here, u_C is the stationary point in the active-CLF region, Π is interval projection, S is the candidate set, and $u^\ddagger$ is the minimum-cost feasible command.

The certificate is conditional on three explicit assumptions: the initial state belongs to C , the scalar interval is nonempty, and online prediction residuals remain exchangeable with the calibration residuals. On the conformal event, Eq. (17) enforces the HOCBF inequality for the declared uncertainty set. The CLF controls practical convergence but is allowed to relax when convergence and safety conflict.

3. EXPERIMENTAL DESIGN

All values are public, normalized simulation parameters and do not identify a specific platform. The experiment retains only evidence that addresses the claimed contributions: residual prediction and calibration, uncertainty-radius

conservatism, time-domain constraint response, Monte Carlo terminal statistics, and a focused predictor ablation.

The three sets jointly vary initial range and closing speed, initial field-of-view angle, target-maneuver amplitude, sensor-noise level, measurement delay, actuator time constant, command limit, and engagement horizon. Within each set, the Monte Carlo campaign further randomizes maneuver phase, measurement noise, and the initial line-of-sight perturbation. Set C is a near-boundary stress condition, but it remains inside the same declared maneuver and sensor generator and is therefore not labeled as a formal out-of-distribution test. The robustness and generalization claims are restricted to this parameterized simulation domain; a deployment shift requires new calibration and independent validation rather than extrapolation of the present coverage result.

Table 1: Learning, calibration, and simulation settings

Item	Value	Role
Total generated samples	87,560	Public normalized maneuver data
Training subset	63,043	Fit the residual output layer
Calibration subset	15,760	Compute the split-conformal rank
Test subset	8,757	Report RMSE and empirical coverage
Hidden units	56	Compact nonlinear basis
Ridge coefficient	4×10^{-4}	Numerical regularization
Miscoverage level	0.05	Nominal 95% marginal coverage
Conformal radius	12.194 m/s ²	Finite-sample calibrated bound
Fixed comparison radius	24.966 m/s ²	99.9% empirical comparison bound
Integration step	0.04 s	Closed-loop propagation
Monte Carlo trials	220 per method/set	Statistical evaluation

Table 2: Evaluation sets from mild to near-boundary severe conditions

Set	A—mild	B—moderate	C—severe
R_0 (m)	6500	8200	10000
V_c (m/s)	500	600	700
q_0 (deg)	8	25	29
A_r (m/s ²)	22	40	55
Noise (m/s ²)	0.8	1.8	3.0
Delay (s)	0.08	0.12	0.16
τ_a (s)	0.06	0.10	0.08
$u_{\max}$ (m/s ²)	115	220	350
T (s)	13.0	13.7	14.3

The reported performance measures are defined as:

$$\begin{aligned} e_f &= R_f |q_f|, \quad J_u = T_s \sum_{k=0}^{N-1} u_k^2, \\ \eta_{\text{FOV}} &= 100 \max_k \frac{|q_k|}{q_{\max}}, \quad \eta_{\text{int}} = \frac{100}{N} \sum_{k=0}^{N-1} \mathbf{1}(|u_k - u_{L,k}| > \varepsilon_u) \end{aligned} \quad (27)$$

Here, e_f is terminal angular cross-track error, R_f and q_f are final range and line-of-sight angle, J_u is command energy, N is the number of samples, η_{FOV} is maximum field-of-view utilization, η_{int} is filter-intervention rate, ε_u is the intervention threshold, and $\mathbf{1}(\cdot)$ is the indicator function.

Table 3: Comparison methods and the role of each component

Method	Learning	Uncertainty	Safety filter	Purpose
Physics-only	No	No	No	Model-based reference
Learning-only	Yes	No	No	Isolate residual-learning value
Fixed-bound CBF	Yes	Fixed 99.9% bound	HOCBF	Compare manual and calibrated bounds
Proposed method	Yes	95% split conformal	CLF-HOCBF	Complete architecture

Each method-set pair uses 220 Monte Carlo trials. Quantitative metrics are computed from raw signals. Savitzky–Golay smoothing is applied only to the displayed time histories so that the visual comparison emphasizes closed-loop trends rather than sample-level noise.

For reproducibility, the learning data are generated from 220 independent 16 s maneuver sequences with a 0.04 s step. For each sequence, the primary and secondary sinusoidal amplitudes are sampled uniformly from 8–55 m/s² and 0–25 m/s², their frequencies from 0.25–1.2 rad/s and 1.2–2.8 rad/s, the phase from $-\pi$ to π , and the transition time from 4–11 s; step and chirp components are each activated with probability 0.5. A 0.12 s measurement delay and Gaussian noise with standard deviation 1.8 m/s² are applied. Samples from indices 4 onward yield 87,560 feature-target rows, which are shuffled with Twister seed 20260720 and divided 72%/18%/10% into training, calibration, and test subsets. The predictor uses 56 tanh random features, hidden weights drawn from $N(0, 0.65^2)$, biases drawn uniformly from -1 to 1, and ridge coefficient 4×10^{-4} . For Monte Carlo trial j , the seed is $100000 + 10000(\text{set index}) + 1000(\text{method index}) + j$; the initial line-of-sight angle receives a 0.5% Gaussian perturbation and the initial angular rate is drawn from $N(0, 0.0014^2)$. At each time step, the implementation filters the delayed measurement, forms the seven-element history feature, predicts the residual maneuver, constructs the nominal command, evaluates the robust HOCBF and softened CLF coefficients, selects the minimum-cost command from the feasible scalar interval and finite candidate set, propagates the first-order actuator and relative dynamics, and stores unsmoothed signals for all quantitative metrics.

4. RESULTS AND DISCUSSION

4.1. Prediction and Calibration

The residual model achieved a held-out RMSE of 5.768 m/s². The empirical coverage was 94.83%, close to the declared 95% marginal target. The conformal radius was 12.194 m/s², whereas the fixed 99.9% comparison radius was 24.966 m/s². Thus, data calibration reduced the uncertainty radius by 51.16% without introducing a Gaussian residual assumption.




Fig. 1. Representative Set B target-maneuver prediction and 95% split-conformal band

Figure 1 shows that the predictor follows the dominant maneuver trend despite delay and measurement noise. The band is not interpreted as a pointwise conditional guarantee; its role is to provide the scalar residual bound used in Eq. (17).




Fig. 2. Set B line-of-sight response of the four comparison methods

4.2. Time-Domain Constraint Response

In Set B, the four trajectories remain inside the $\pm 30^\circ$ field-of-view boundary, while the learning-enhanced responses reduce the late-stage residual relative to physics-only guidance. The safety-filtered responses remain close to the learning-only command because the optimization applies the smallest admissible correction.




Fig. 3. Field-of-view utilization of the proposed method for Sets A, B, and C

The three-set comparison makes the operating-domain progression explicit. Set C begins close to the field-of-view boundary and therefore tests constraint handling rather than nominal accuracy. The time histories remain below the declared limit and converge toward the terminal region without a separate barrier-margin plot.

The correction is concentrated in the initial high-utilization interval and then decays as the line of sight moves away from the boundary. This is consistent with the median intervention rates of 65.49% in Set B and 54.93% in Set C: the filter is materially active, but it does not replace the nominal guidance law throughout the engagement.




Fig. 4. Normalized time-varying correction applied by the solver-free safety filter in the three sets




Fig. 5. Normalized inner-loop command-tracking error in the three sets

The actuator error is largest during the initial command transient and at maneuver reversals. Its explicit inclusion in Eqs. (17) and (21) prevents the safety argument from assuming instantaneous acceleration realization.

4.3. Monte Carlo results

This is consistent with the median intervention rates of 65.49% in Set B and 54.93% in Set C: the filter is materially active, but it does not replace the nominal guidance law throughout the engagement.

Table 4: Monte Carlo results

Set	Method	Median e_t (m)	P95 e_t (m)	Max. FOV (%)	Filter active (%)
A	Physics-only	0.0889	0.1972	26.68	0.00
A	Learning-only	0.1040	0.1635	26.67	0.00
A	Fixed-bound CBF	0.1075	0.1627	26.66	24.61
A	Proposed method	0.1179	0.1616	26.66	24.30
B	Physics-only	0.1420	0.2760	83.30	0.00
B	Learning-only	0.0932	0.1756	83.28	0.00
B	Fixed-bound CBF	0.1043	0.1730	83.39	65.49
B	Proposed method	0.0957	0.1787	83.34	65.49
C	Physics-only	0.3676	0.4556	96.58	0.00
C	Learning-only	0.3582	0.4196	96.73	0.00
C	Fixed-bound CBF	0.3592	0.4179	96.68	54.93
C	Proposed method	0.3607	0.4164	96.64	54.93

Set B provides the clearest performance evidence. Relative to physics-only guidance, the proposed method reduces median terminal error from 0.1420 m to 0.0957 m, a 32.64% reduction, and reduces P95 from 0.2760 m to 0.1787 m, a 35.26% reduction. The learning-only median is slightly lower, but it has no calibrated uncertainty certificate. The fixed-bound method has similar tail performance, yet its uncertainty radius is more than twice the conformal radius.

Set C is safety dominated. The proposed method reduces P95 by 8.60% relative to physics-only guidance while operating at 96.64% median maximum field-of-view utilization. Its median error differs little from the other learning-based variants because command saturation and the near-boundary constraint dominate the central tendency. In Set A, physics-only has the smallest median, whereas the proposed method has the smallest P95. This trade-off is retained because it defines the method's actual value: improved tail robustness and an auditable uncertainty interface, not universal nominal superiority.

These results define the operating conditions under which the proposed architecture is meaningful. In a mild engagement such as Set A, the physics model is already adequate and the additional learned feedforward and safety-filter actions can slightly increase the median error; the proposed method is

therefore not recommended solely to optimize nominal median accuracy in that regime. Its clearest value arises when maneuver uncertainty, measurement delay, and noise affect the upper tail of terminal error, or when the line of sight approaches the field-of-view boundary. In those conditions, method selection should emphasize tail risk, calibrated uncertainty traceability, and constraint handling rather than the median metric alone.

No field-of-view violations were observed in the declared Monte Carlo domain for any method. Accordingly, the result is not presented as an empirical violation-rate reduction. The safety contribution is the explicit robust inequality, the nonempty scalar projection, and measurable filter activity near the boundary.

4.4. Focused Predictor Ablation

Replacing the residual predictor with a raw delayed estimate increases the median error by 37.9% in Set B and 3.2% in Set C. The larger Set B effect confirms that the predictor adds value when delay and maneuver are substantial but the engagement is not yet dominated by hard saturation. The smaller Set C difference is physically consistent: once the barrier and command limit dominate, estimator improvements have less authority to change the final response.

Table 5 presents a focused ablation study conducted separately from the principal Monte Carlo campaign, using 90 trials per variant and a distinct random-seed sequence. Consequently, the median values reported in Table 5 are not expected to coincide exactly with the corresponding medians obtained from the 220-trial Monte Carlo simulations in Table 4.

The comparison methods also provide a nested component-level ablation. Physics-only versus learning-only isolates residual learning; learning-only versus fixed-bound CBF shows the effect of activating the safety filter; and fixed-bound CBF versus the proposed method isolates replacement of the fixed uncertainty envelope by the split-conformal radius. The raw-delay variant further isolates the history-based predictor. CLF relaxation is not disabled as a separate variant because removing it can make the convergence inequality incompatible with the hard HOCBF and actuator limits; its practical role is instead monitored through filter intervention, interval feasibility, and infeasible-step counts. Sensitivity to calibration design follows the order statistic in Eq. (11): increasing the requested confidence cannot decrease the calibrated radius, whereas reducing the calibration-set size makes the attainable quantile levels coarser. A full numerical sweep over calibration size, confidence level, conformal radius, and CLF penalty is not reported here and remains necessary before drawing broader tuning conclusions.

Table 5: Focused ablation of the residual predictor

Set	Variant	Median e_t (m)	Change from full method
B	Full proposal	0.0985	—
B	Raw delayed estimate	0.1358	+37.9%
C	Full proposal	0.3549	—
C	Raw delayed estimate	0.3662	+3.2%

5. CONCLUSION

This study has developed a certifiable learning-enhanced



guidance law by combining physics-based nominal guidance, residual maneuver prediction, split-conformal uncertainty calibration, and robust CLF–HOCBF safety filtering. The main achievement is that the learning error is not treated as an exact estimate; instead, it is converted into a finite-sample uncertainty radius and inserted directly into the field-of-view constraint. The learned module therefore improves maneuver compensation, while constraint enforcement remains the responsibility of an explicit safety filter.

Future work should investigate online or weighted conformal recalibration under distribution shift, extension to full three-dimensional and six-degree-of-freedom dynamics, additional command-rate and asymmetric-seeker constraints, hardware-in-the-loop timing assessment, and validation using semi-physical or experimental data. These developments are required to evaluate the certificate when exchangeability and inner-loop-model assumptions are only approximately satisfied.

DECLARATIONS

Author Contributions

Trinh Thi Minh and Vu Quang Luong jointly contributed to the conceptualization, methodology, simulation design, analysis, validation, and manuscript preparation.

Vu Quang Luong also served as the corresponding author. Both authors reviewed and approved the final manuscript.

Conflict of Interest

The authors declare that they have no conflict of interest.

Funding

This research received no external funding.

Institutional Clearance for Open-Access Publication

The authors confirm that this work has been cleared by the Air Defence-Air Force Academy for open-access publication. The study uses only public, normalized simulation parameters and contains no classified, restricted, or platform-identifying information.

REFERENCES

- [1] X. Tang, J. Yu, X. Dong, and Z. Ren, “Integrated guidance and control with impact angle and general field-of-view constraints,” *Aerospace Science and Technology*, vol. 144, Art. no. 108809, 2024, doi: [10.1016/j.ast.2023.108809](https://doi.org/10.1016/j.ast.2023.108809).
- [2] B. Gaudet, R. Furfaro, and R. Linares, “Reinforcement learning for angle-only intercept guidance of maneuvering targets,” *Aerospace Science and Technology*, vol. 99, Art. no. 105746, 2020, doi: [10.1016/j.ast.2020.105746](https://doi.org/10.1016/j.ast.2020.105746).
- [3] X. Gong, W. Chen, and Z. Chen, “All-aspect attack guidance law for agile missiles based on deep reinforcement learning,” *Aerospace Science and Technology*, vol. 127, Art. no. 107677, 2022, doi: [10.1016/j.ast.2022.107677](https://doi.org/10.1016/j.ast.2022.107677).
- [4] N. Wang, X. Wang, N. Cui, Y. Li, and B. Liu, “Deep reinforcement learning-based impact time control guidance law with constraints on the field-of-view,” *Aerospace Science and Technology*, vol. 128, Art. no. 107765, 2022, doi: [10.1016/j.ast.2022.107765](https://doi.org/10.1016/j.ast.2022.107765).
- [5] W. Chen, C. Gao, and W. Jing, “Proximal policy optimization guidance algorithm for intercepting near-space maneuvering targets,” *Aerospace Science and Technology*, vol. 132, Art. no. 108031, 2023, doi: [10.1016/j.ast.2022.108031](https://doi.org/10.1016/j.ast.2022.108031).
- [6] Z. Yang, X. Liu, and H. Liu, “Impact time control guidance law with time-varying velocity based on deep reinforcement learning,” *Aerospace Science and Technology*, vol. 142, Art. no. 108603, 2023, doi: [10.1016/j.ast.2023.108603](https://doi.org/10.1016/j.ast.2023.108603).
- [7] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017, doi: [10.1109/TAC.2016.2638961](https://doi.org/10.1109/TAC.2016.2638961).
- [8] W. Xiao and C. Belta, “High-order control barrier functions,” *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3655–3662, 2022, doi: [10.1109/TAC.2021.3105491](https://doi.org/10.1109/TAC.2021.3105491).
- [9] A. J. Taylor and A. D. Ames, “Adaptive safety with control barrier functions,” in *Proceedings of the American Control Conference*, Denver, USA, 2020, pp. 1399–1405, doi: [10.23919/ACC45564.2020.9147463](https://doi.org/10.23919/ACC45564.2020.9147463).
- [10] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, “End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 1, pp. 3387–3395, 2019, doi: [10.1609/aaai.v33i01.33013387](https://doi.org/10.1609/aaai.v33i01.33013387).
- [11] Z. Marvi and B. Kiumarsi, “Safe reinforcement learning: A control barrier function optimization approach,” *International Journal of Robust and Nonlinear Control*, vol. 31, no. 6, pp. 1923–1940, 2021, doi: [10.1002/mc.5132](https://doi.org/10.1002/mc.5132).
- [12] J. Lei, M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman, “Distribution-free predictive inference for regression,” *Journal of the American Statistical Association*, vol. 113, no. 523, pp. 1094–1111, 2018, doi: [10.1080/01621459.2017.1307116](https://doi.org/10.1080/01621459.2017.1307116).
- [13] A. N. Angelopoulos and S. Bates, “Conformal prediction: A gentle introduction,” *Foundations and Trends in Machine Learning*, vol. 16, no. 4, pp. 494–591, 2023, doi: [10.1561/22000000101](https://doi.org/10.1561/22000000101).