

ROBUST ESTIMATION-AWARE GUIDANCE FOR STOCHASTIC INTERCEPTION WITH PARAMETRIC UNCERTAINTY

Liraz Mudrik[1] & Yaakov Oshman[2]

[1]Naval Postgraduate School, Monterey, CA, 93943
[2]Technion—Israel Institute of Technology, Haifa, 3200003, Israel

Abstract

When a target's agility and maneuverability are unknown, the boundary between capture and miss becomes blurred, causing standard interceptors to fail. This work challenges the traditional separation of estimation and control by proposing a robust guidance strategy that explicitly accounts for this uncertainty. Instead of unrealistically assuming perfect knowledge of the target's parameters, our approach utilizes a risk-aware decision mechanism that weighs the consequences of multiple target hypotheses. Furthermore, we introduce a trajectory shaping strategy: when the optimal control is nonunique, the interceptor adapts its trajectory to maximize information gain. Simulations demonstrate that this unified approach prevents the catastrophic performance drop-off seen in classical methods, restoring high kill probabilities against threats with uncertain dynamic parameters.

Keywords: Estimation-Aware Guidance, Missile Guidance, Pursuit-Evasion Differential Games, Robust Control, Trajectory Shaping

1 Introduction

The interception of highly maneuvering targets, such as tactical ballistic missiles, is a stochastic optimal control problem with imperfect information. Typically invoking the separation principle, the classical approach consists of a separate and independent design of the estimator and the guidance law. While optimal for linear-Gaussian systems, this separation is suboptimal for stochastic systems involving nonlinearities and non-Gaussian distributions. In such cases, the generalized separation theorem (GST) should be obeyed [1]. The GST states that the estimator can be designed independently of the guidance law, but the control law must take into account the entire posterior distribution provided by the estimator. Despite this, plugging point estimates into a guidance law derived under perfect information remains the standard practice, and is known to yield substantial performance loss in realistic scenarios [2].

Several stochastic guidance laws addressing this issue have been proposed in the literature. Speyer [3] introduced an adaptive terminal scheme based on an exponential cost criterion. The works [4] and [5] developed nonlinear stochastic optimal control and LQG-based guidance laws with bounded acceleration commands. An integrated estimation-guidance design approach against randomly maneuvering targets was proposed in [6]. Later, [7] used the notion of reachability sets to develop a GST-compliant methodology that does not rely on linearity or Gaussianity assumptions.

The baseline guidance law considered in this work is the differential game-based guidance law DGL1, derived as the saddle-point solution of a linearized zero-sum differential game between an interceptor and a target with first-order dynamics [8]. Under the perfect-information assumption, an interceptor with superior maneuverability and agility is guaranteed hit-to-kill performance, with a singular region

forming in the game space whose boundary depends on the time-to-go and on the maximal acceleration capabilities and time constants of both players [9]. DGL1's strategy depends on the relative position of the game's state with respect to this boundary.

In a recent work [10], the authors introduced a unified estimation-guidance framework that modifies DGL1 to comply with the GST through a multihypothesis Bayesian decision algorithm. The resulting guidance law, which we shall henceforth refer to as estimation-aware guidance (EAGL), uses an interacting multiple model particle filter (IMMPF) [11] to obtain the posterior probability density function of the game's state, and selects the guidance command that minimizes the Bayesian risk. When the resulting decision is nonunique, this ambiguity is exploited via information-enhancement trajectory shaping (IETS) to maximize the information gained by the estimator about the target's state. The framework of [10], however, assumes perfect knowledge of the target's dynamic parameters, namely its time constant τ_T and maximal lateral acceleration $a_T^{\max}$. Also commonplace throughout the broader stochastic guidance literature [12, 4, 5, 7], this assumption is unrealistic. Real-world tactical ballistic missile defense scenarios involve substantial uncertainty regarding target agility, because the interceptor only knows the family of threats to which the target belongs and not the specific airframe within that family. When these parameters are uncertain, the singular region boundary becomes uncertain too, and the interceptor cannot reliably evaluate the state of the game. Furthermore, a smart target can exploit this uncertainty to enforce a substantial miss distance through a well-timed evasion maneuver [13, 12].

Robust counterparts of DGL-type laws have been studied in the deterministic setting. In [14], the linear pursuit-evasion games were extended to time-varying parameters and bounded controls. Then, [15] derived a guaranteed-miss formulation for systems with bounded controls and bounded noise. These works demonstrate the feasibility of robust DGL-type designs in deterministic settings, but they do not address the stochastic-information case in which the state itself must be estimated from noisy measurements.

The main contribution of this paper is the extension of the unified estimation-guidance framework of [10] to the case of uncertain target dynamic parameters. Two new guidance laws are developed in the ensuing. The first is the robust estimation-aware guidance law (REAGL), in which the Bayesian decision criterion explicitly accounts for the joint posterior distribution of the game's state and the target's parameters. Each particle in the IMMPF is associated with its own dynamic model, and the Bayesian risk is computed against a particle-specific singular region boundary. The second is the robust information-enhancement trajectory shaping law (RIETS), which exploits the nonuniqueness of the optimal Bayesian decision to sharpen the state estimate in the directions most relevant to the parameter-dependent guidance command, by minimizing the determinant of the submatrix of the posterior inverse Fisher information matrix (FIM) tied to the zero-effort miss and the time-to-go. To the best of the authors' knowledge, the present work is the first to address parametric uncertainty in DGL-type guidance via Bayesian decision theory and information-enhancing trajectory shaping in a unified, GST-compliant manner. A Monte Carlo simulation study involving a tactical ballistic missile defense scenario demonstrates that the proposed strategies prevent the catastrophic performance degradation observed when applying perfect-information laws under parametric uncertainty, and recover much of the performance attainable under perfect parameter knowledge.

The remainder of this paper is organized as follows. Section 2 formulates the stochastic interception problem with parametric uncertainty. Section 3 provides a brief, self-contained review of the estimation-aware guidance framework of [10] that constitutes the starting point of the present work. Section 4 presents the proposed REAGL and RIETS laws. A simulation study and a comparative performance analysis are presented in Sec. 5, followed by concluding remarks in Sec. 6.

2 Problem Formulation

We consider a planar, single-pursuer, single-target endgame engagement, as shown in Fig. 1, where X_I-O_I-Y_I is a Cartesian inertial reference frame. The interceptor and the target, as well as variables associated with them, are denoted by M and T, respectively. The speed, lateral acceleration, and path angle are denoted by V, a, and γ, respectively. The slant range between the interceptor and the target is ρ, and the line-of-sight (LOS) angle, measured with respect to the X_I axis, is λ. The

engagement is modeled as a zero-sum differential game in which the missile minimizes, and the target maximizes, the terminal miss distance.

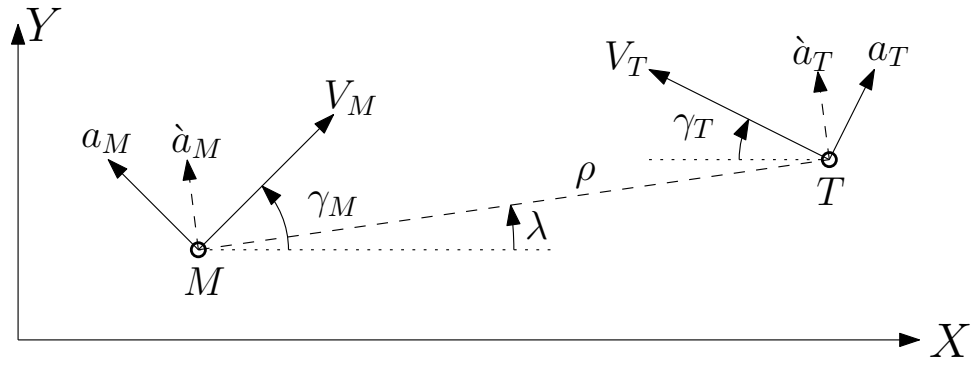


Figure 1 – Planar engagement geometry.

The following underlying assumptions are made, the first five of which paralleling those of [10]:

1) Both the interceptor and the target are represented as point masses.

2) The interceptor's own path angle and lateral acceleration are known (e.g., from its onboard navigation system).

3) The speeds of the interceptor and the target, V_M and V_T, respectively, are known and time-invariant.

4) The target's acceleration command, u_T, belongs to a known admissible set comprising two bang-bang maneuvers,

$$u_T = \begin{cases} +a_T^{\max} & n = 1 \\ -a_T^{\max} & n = 2 \end{cases} \tag{1}$$

where n denotes the target maneuver mode. This is known to be the optimal, perfect-information evasion strategy in bounded-acceleration pursuit-evasion games [8, 9].

5) The interceptor and the target possess first-order acceleration response dynamics, with time constants τ_M and τ_T, respectively. The interceptor's lateral acceleration command, u_M, is bounded in magnitude by $a_M^{\max}$, and its parameters τ_M and $a_M^{\max}$ are known constants.

6) In contradistinction to [10], the target's parameters τ_T and $a_T^{\max}$ are unknown, taking values in known intervals:

$$\tau_T \in [\tau_{T,1}, \tau_{T,2}], \tag{2}$$

$$a_T^{\max} \in [a_{T,1}^{\max}, a_{T,2}^{\max}]. \tag{3}$$

Following common practice [9, 10, 12], we use polar coordinates to describe the nonlinear engagement kinematics. The natural state vector of the engagement,

$$[\rho,\ \lambda,\ \gamma_T,\ a_T]^T, \tag{4}$$

evolves according to the equations of motion

$$\dot{\rho} = V_\rho, \tag{5a}$$

$$\dot{\lambda} = V_\lambda / \rho, \tag{5b}$$

$$\dot{\gamma}_T = a_T / V_T, \tag{5c}$$

$$\dot{a}_T = -\frac{a_T}{\tau_T} + \frac{u_T}{\tau_T}, \tag{5d}$$

where the relative velocity components along and perpendicular to the LOS, respectively, are

$$V_\rho = -(V_M \cos\delta_M + V_T \cos\delta_T), \tag{6a}$$

and

$$V_\lambda = -V_M \sin\delta_M + V_T \sin\delta_T, \tag{6b}$$

and the lead angles are $\delta_M = \gamma_M - \lambda$ and $\delta_T = \gamma_T + \lambda$. By Assumption 2), the interceptor's own path angle and lateral acceleration evolve according to the known equations

$$\dot{\gamma}_M = a_M / V_M, \tag{7a}$$

$$\dot{a}_M = -\frac{a_M}{\tau_M} + \frac{u_M}{\tau_M}, \tag{7b}$$

and are, therefore, available to the estimator at all times.

During the short endgame phase, and provided the engagement remains near the initial collision course, the standard linearization of the kinematics with respect to the initial LOS holds [9, 12]. Under this approximation, the time-to-go until intercept can be expressed as $t_{\text{go}} = -\rho/V_\rho$, and the zero-effort miss (ZEM), z_1, namely the predicted miss distance assuming that no further controls are applied by either player, is given by [9]

$$z_1(t) = z_{\text{PN}}(t) - a_M \tau_M^2\, \psi(\eta_M) + a_T \tau_T^2\, \psi(\eta_T), \tag{8}$$

where z_{PN} is the proportional-navigation-style ZEM, $\eta_i \triangleq t_{\text{go}}/\tau_i$ for $i \in \{M,T\}$, and $\psi(\mu) \triangleq e^{-\mu} + \mu - 1$. Following [10], we define the state vector used by both the estimator and the guidance law as

$$\mathbf{x}_M = \begin{bmatrix} t_{\text{go}} & z_1 & \lambda & \gamma_T \end{bmatrix}^T. \tag{9}$$

The choice of this state vector, rather than the natural one in (4), is motivated by the fact that the DGL1 game space, and, in particular, the singular region boundary introduced in Section 3, is parameterized by exactly t_{go} and z_1. The nonlinear equations of motion governing (9), $\dot{\mathbf{x}}_M = f(\mathbf{x}_M, u_M, u_T; \theta_T)$, are derived in [10] from (5)–(6); we omit them here for conciseness.

The interceptor uses an angle-only sensor that provides line-of-sight angle measurements corrupted by zero-mean Gaussian noise of standard deviation σ,

$$y^k = \lambda^k + \nu^k, \qquad \nu^k \sim \mathcal{N}(0, \sigma^2), \tag{10}$$

sampled at frequency f. We let $\mathscr{Y}^k \triangleq \{y^1, \ldots, y^k\}$ denote the measurement history available at time t^k.

Assumption 6 constitutes the central technical departure of this work from the framework of [10], in which both target parameters are taken to be known. We collect the uncertain target parameters as $\theta_T \triangleq [\tau_T, a_T^{\max}]^T$, taking values in $\Theta_T \triangleq [\tau_{T,1}, \tau_{T,2}] \times [a_{T,1}^{\max}, a_{T,2}^{\max}]$. Because the singular region boundary of the underlying differential game, defined in the next section, depends on both components of θ_T, the interceptor cannot deterministically evaluate the state of the game, nor can it determine which guidance command is optimal, without explicitly reasoning over θ_T. This is the operational obstacle that the proposed robust framework addresses.

3 Estimation-Aware Guidance Framework: A Brief Review

This section briefly summarizes the estimation-aware guidance framework of [10], which serves as the starting point for the robust extensions developed in this paper. For detailed derivations and proofs, the reader is referred to [10].

3.1 DGL1 Game Space

Under the perfect-information assumption, that is, with θ_T known and the state of the game directly observable, the game space of the linearized zero-sum interception game can be partitioned into a regular region $\mathscr{D}_1$, where bang-bang strategies are optimal, and a singular region $\mathscr{D}_0$, where the interceptor's optimal command is any admissible command [9]. The regular region consists of an upper component $\mathscr{D}_1^+$, above the singular region, and a lower component $\mathscr{D}_1^-$, below it, as shown in Fig. 2. The boundary of the singular region, denoted $z_1^*(t_{\text{go}})$, depends on t_{go} and on the players' parameters; an explicit expression is given in [9, 10]. Throughout the present section, the target parameter θ_T is a known constant, and is therefore suppressed from the boundary notation; its explicit dependence is reinstated in Section 4, where θ_T becomes uncertain. The DGL1 control law issues a maximum-magnitude command of appropriate sign in $\mathscr{D}_1$ and a saturated linear command in $\mathscr{D}_0$, the latter being used to prevent control chattering. The saturated linear command applied inside $\mathscr{D}_0$ is

$$u_M^k = a_M^{\max}\,\mathrm{sat}\left(\frac{z_1(t_{\text{go}})}{k\,z_1^*(t_{\text{go}})}\right), \qquad 0 < k \leq 1, \tag{11}$$

where $\mathrm{sat}(\cdot)$ is the saturation function and k sets the fraction of the singular region over which the command is linear.

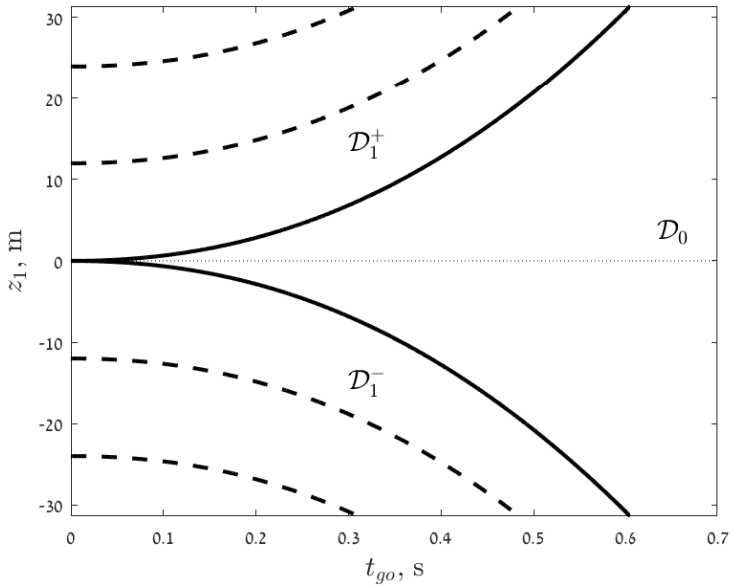


Figure 2 – DGL1 game space, partitioned into the regular regions $\mathscr{D}_1^-$ and $\mathscr{D}_1^+$, and the singular region $\mathscr{D}_0$. The solid lines are the boundaries $\pm z_1^*(t_{\text{go}})$ of the singular region.

3.2 Interacting Multiple Model Particle Filter

The interceptor estimates the joint posterior of the game state and the target maneuver mode using an interacting multiple model particle filter (IMMPF) [11]. The choice of an IMMPF is motivated by three properties of the problem at hand. First, the nonlinear EOMs governing (9), derived from (5)–(6), preclude closed-form filtering. Second, the target's bang-bang acceleration command (1) renders the system hybrid, with a continuous state and a discrete maneuver mode that must be jointly estimated. Third, the propagation of the line-of-sight measurement (10) through the polar-to-Cartesian kinematics produces a non-Gaussian posterior, even when the measurement noise itself is Gaussian. At each discrete time t^k, the IMMPF approximates the joint posterior $p(\mathbf{x}^k, n^k \mid \mathscr{Y}^k)$ by a weighted particle ensemble

$$\{(\mathbf{x}_{j'}^k, n^{j'}, w_{j'}^k) : j' = 1, \ldots, N_p\}, \tag{12}$$

where $\mathbf{x}_{j'}^k$ is the continuous state of the j'th particle, $n^{j'}$ is its associated discrete mode, $w_{j'}^k$ is its weight, and N_p is the total particle population. The weights satisfy $\sum_{j'=1}^{N_p} w_{j'}^k = 1$, and the mode-

conditional posterior is

$$\Pr(n^k = n \mid \mathscr{Y}^k) \approx \sum_{j':n^{j'}=n} w^k_{j'}. \tag{13}$$

The target's discrete maneuver mode is modeled as a non-homogeneous Markov chain with time-varying transition probabilities,

$$\pi^k_{n,n'} \triangleq \Pr(n^k = n' \mid n^{k-1} = n). \tag{14}$$

The non-homogeneity captures the behavior of a smart target [13]: such a target times its acceleration switch within a specific window of the engagement to maximize the resulting miss distance. A homogeneous chain would implicitly assume that the evader is indifferent to the timing of the switch and would, in turn, fail to capture this aspect of optimal evasion. The interested reader is referred to [11] for a complete derivation and implementation details of the IMMPF.

3.3 Bayesian Decision-Based Guidance

In the stochastic case, the state of the game is not deterministically known. To address this, the four hypotheses defined in [10] categorize the possible states of the game:

H_1**:** the state is in the upper regular region.

H_2**:** the state is in $\mathscr{D}_0$, with target maneuver mode $n = 1$.

H_3**:** the state is in $\mathscr{D}_0$, with target maneuver mode $n = 2$.

H_4**:** the state is in the lower regular region.

These four hypotheses are illustrated in Fig. 3, which shows a particle cloud, representing the posterior distribution over the game state, overlaid on the game space partition.

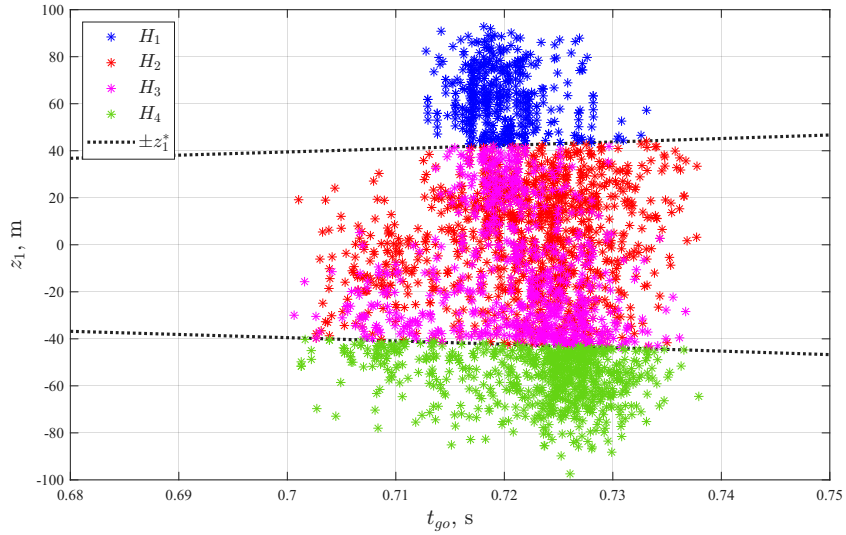


Figure 3 – IMMPF particle cloud at a particular moment during the engagement. The particles are color-coded according to the hypotheses to which they belong. The black dotted lines are the boundaries of the singular region.

The framework formulates the stochastic guidance problem as a multi-hypothesis Bayesian decision problem [16]. Let C_{ij} denote the cost (final miss distance) incurred when applying the guidance command associated with H_i when, in fact, H_j is true. Letting m denote the number of hypotheses (in [10] $m = 4$) and P_j denote the prior probability of hypothesis H_j, the unnormalized additional risk of deciding hypothesis H_i is [10]

$$I_i(\mathscr{Y}) \triangleq \sum_{\substack{j=1 \\ j\neq i}}^{m} P_j \Pr(\mathscr{Y} \mid H_j)(C_{ij} - C_{jj}), \quad i = 1,\ldots,m \tag{15}$$

where $\Pr(\mathscr{Y} \mid H_j)$ is the likelihood of observing the measurements given that hypothesis H_j is true. The optimal Bayesian decision rule is

$$\text{Decide } H_{i^\star}, \quad i^\star \triangleq \arg\min_{i\in\{1,\dots,m\}} I_i(\mathscr{Y}). \tag{16}$$

Under the assumption that wrong decisions are no cheaper than correct ones, that is, $C_{ij} \geq C_{jj}$ for all i, j, this rule minimizes the decision-dependent component of the total Bayesian risk.

Implementing the decision rule (16) at each time t^k requires the computation of three components per hypothesis: the likelihoods $\{\Pr(\mathscr{Y} \mid H_j)\}_{j=1}^m$, the costs $[C_{ij}]_{i,j=1}^m$, and the prior probabilities $\{P_j\}_{j=1}^m$. All three are obtained from the IMMPF particle ensemble (12). We describe each in turn.

We begin with the likelihoods. We approximate the likelihood of each hypothesis H_j by summing the IMMPF weights of the particles associated with it, i.e., the particles whose state and mode values place them in H_j,

$$\Pr(\mathscr{Y} \mid H_j) \approx \sum_{j'\in H_j} w_{j'}^k, \tag{17}$$

where the notation $j' \in H_j$ stands for "particle j' whose state $(z_1^{j'}, t_{\mathsf{go}}^{j'})$ and mode $n^{j'}$ satisfy the indicator condition defining H_j." For example, $j' \in H_2$ requires both that $|z_1^{j'}| \leq z_1^*(t_{\mathsf{go}}^{j'})$ and that $n^{j'} = 1$.

Turning next to the costs, we recall that C_{ij} represents the expected terminal miss distance accumulated when the interceptor applies the guidance command associated with hypothesis H_i while H_j is, in fact, the correct hypothesis. We define the normalized weight of particle j' within hypothesis H_j as

$$\tilde{w}_{j'}^k \triangleq \frac{w_{j'}^k}{\sum_{j''\in H_j} w_{j''}^k}, \quad j' \in H_j, \tag{18}$$

which is well-defined whenever $\Pr(\mathscr{Y} \mid H_j) > 0$. For the cost of a correct decision, that is, when $i = j$, the IMMPF-based approximation reads [10]

$$C_{jj} = \sum_{j'\in H_j} \tilde{w}_{j'}^k \max\left(0,\ \left|z_1^{j'}(t_{\mathsf{go}}^{j'})\right| - z_1^*(t_{\mathsf{go}}^{j'})\right), \tag{19}$$

where $z_1^{j'}$ and $t_{\mathsf{go}}^{j'}$ are the ZEM and the time-to-go associated with particle j', and $z_1^*(t_{\mathsf{go}}^{j'})$ is the singular boundary at that time-to-go. In the case of an incorrect decision, that is, when $i \neq j$, when both $\Pr(\mathscr{Y} \mid H_i) > 0$ and $\Pr(\mathscr{Y} \mid H_j) > 0$, the cost is [10]

$$C_{ij} = \sum_{j'\in H_j} \tilde{w}_{j'}^k \sum_{i'\in H_i} \tilde{w}_{i'}^k \max\left(0,\ \left|z_1^{i'j'}(t_{\mathsf{go}}^{j'} - h)\right| - z_1^*(t_{\mathsf{go}}^{j'} - h)\right), \tag{20}$$

where $z_1^{i'j'}(t_{\mathsf{go}}^{j'} - h)$ denotes the ZEM of particle $j' \in H_j$ propagated forward over the interval $[t_{\mathsf{go}}^{j'}, t_{\mathsf{go}}^{j'} - h]$ while applying the acceleration command of particle $i' \in H_i$, assumed constant over the interval. The parameter $h > 0$ is a predefined cost-evaluation horizon. In the two remaining, degenerate, cases, the cost expression simplifies. When $\Pr(\mathscr{Y} \mid H_i) = 0$ while $\Pr(\mathscr{Y} \mid H_j) > 0$, no particle is available in H_i, and the inner sum over $i' \in H_i$ in (20) collapses, leaving the single sum

$$C_{ij} = \sum_{j'\in H_j} \tilde{w}_{j'}^k \max\left(0,\ \left|z_1^{i'j'}(t_{\mathsf{go}}^{j'} - h)\right| - z_1^*(t_{\mathsf{go}}^{j'} - h)\right), \tag{21}$$

where the superscript i' now refers to the command of hypothesis H_i rather than to a populated particle. This command is $+a_M^{\max}$ for H_1, $-a_M^{\max}$ for H_4, and the saturated linear command for H_2 and H_3. When $\Pr(\mathscr{Y} \mid H_j) = 0$, hypothesis H_j carries zero posterior weight and drops out of the risk (15), so C_{jj} and C_{ij} may be assigned arbitrarily, consistent with line 14 of Algorithm 1.

Finally, we turn to the prior probabilities. The prior probability P_j of hypothesis H_j at time t^k, conditioned on the measurement history available at time t^{k-1}, is obtained from the law of total probability over the switching event in $[t^{k-1}, t^k]$,

$$P_j = \Pr(H_j \mid \mathscr{Y}^{k-1}, \mathsf{SW})\ \Pr(\mathsf{SW}) + \Pr(H_j \mid \mathscr{Y}^{k-1}, \mathsf{NSW})\ \Pr(\mathsf{NSW}), \tag{22}$$

where SW denotes the event that the target has switched its acceleration command in the interval $[t^{k-1}, t^k]$, and NSW denotes its complement. The conditional hypothesis probabilities $\Pr(H_j \mid \mathscr{Y}^{k-1}, \mathrm{SW})$ and $\Pr(H_j \mid \mathscr{Y}^{k-1}, \mathrm{NSW})$ are obtained by propagating the particle cloud available at time t^{k-1} forward under the SW and NSW scenarios, respectively, and summing the resulting weights of the particles that fall in H_j. The probabilities $\Pr(\mathrm{SW})$ and $\Pr(\mathrm{NSW})$ are determined by the non-homogeneous transition probability matrix (TPM) (14).

The probabilities P_j, the likelihoods $\Pr(\mathscr{Y} \mid H_j)$, and the costs C_{ij} are thus all computed from the particle approximation (12) produced by the IMMPF [11, 17]. Two hypotheses ($n = 1, 2$) for the target acceleration command are used in the nominal framework of [10], an assumption that will be relaxed in Section 4.

Direct computation of I_i is expensive. Lemma 1 of [10] provides necessary and sufficient conditions for $I_i = 0$ in terms of region indicators $\mathscr{A}_i \subset \mathscr{D}_0 \cup \mathscr{D}_1$ defined over the current particle states. When the particle cloud satisfies $\sum_{j \in \mathscr{A}_i} w_j^k = 1$ for some hypothesis i, the decision (16) is optimal and the explicit cost computation is bypassed, yielding a substantial reduction in the computational effort. Algorithm 1 summarizes one cycle of the resulting decision criterion.

Algorithm 1 One cycle of the Bayesian decision criterion, adopted from [10].

Obtain particle weights and samples at times t^{k-1} and t^k.
for $j = 1, \ldots, m$ **do**
  Compute the likelihood probability of H_j via Eq. (17).
  if $\Pr(\mathscr{Y} \mid H_j) > 0$ **then**
    for $i = 1, \ldots, m$ **do**
      if $i \neq j$ **then**
        Compute the cross cost C_{ij} via Eq. (20).
      else
        Compute the diagonal cost C_{jj} via Eq. (19).
      end if
    end for
    Compute the prior probability P_j via Eq. (22).
  else
    Set $C_{ij} \leftarrow 0$ for all i, and $P_j \leftarrow 0$.
  end if
  Compute the expected excess cost I_j via Eq. (15).
end for
Apply the optimal Bayesian decision rule via Eq. (16).

3.4 Information-Enhancement Trajectory Shaping

The Bayesian decision criterion may yield nonunique solutions, that is, $I_i = 0$ for multiple hypotheses, all corresponding to states inside $\mathscr{D}_0$. In this case, the interceptor's command is admissible over a set $\mathscr{U}_M^k$ [10],

$$\mathscr{U}_M^k \triangleq \left\{ u_M^k : \left| u_M^k \right| \leq a_M^{\max}, \sum_{j' \in \{H_1, H_4\}} w_{j'}^{k+1} \leq W_{\mathrm{Thres}} \right\} \tag{23}$$

which keeps most of the particles inside the singular region while allowing some particles outside, with the total weight of those outside bounded by a predetermined value W_{Thres}. IETS exploits this freedom to maximize the information available to the estimator. Specifically, it selects the command so as to minimize the determinant of the submatrix of the inverse FIM corresponding to z_1 and t_{go},

$$u_M^k = \arg \min_{u_M^k \in \mathscr{U}_M^k} \det \left((\bar{J}^{k+\bar{h}})^{-1} \right) \tag{24}$$

where $\bar{h}$ is a predefined trajectory-shaping horizon and $(\bar{J}^{k+\bar{h}})^{-1}$ is the (z_1, t_{go}) subblock of the inverse FIM, namely the predicted posterior CRLB at time $t^{k+\bar{h}}$. The use of $\det J$ as an observability metric

follows [18]. The particle-based approximation of the FIM follows [19], building on the posterior Cramér–Rao lower bound (CRLB) recursion of [20].

3.5 Estimation-Aware Guidance Law

Combining the Bayesian decision criterion with the saturated linear command (11) in the no-decision branch yields the estimation-aware guidance law (EAGL) [10],

$$u_M^k = \begin{cases} +a_M^{\max} & H_1 \text{ decided} \\ \displaystyle\sum_{j' \in H_2} \tilde{w}_{j'}^k \frac{z_1^{j'}(t_{\mathrm{go}}^{j'})}{z_1^*(t_{\mathrm{go}}^{j'})} a_M^{\max} & H_2 \text{ decided} \\ \displaystyle\sum_{j' \in H_3} \tilde{w}_{j'}^k \frac{z_1^{j'}(t_{\mathrm{go}}^{j'})}{z_1^*(t_{\mathrm{go}}^{j'})} a_M^{\max} & H_3 \text{ decided} \\ -a_M^{\max} & H_4 \text{ decided} \\ a_M^{\max} \operatorname{sat}\left(\dfrac{z_1(t_{\mathrm{go}})}{k z_1^*(t_{\mathrm{go}})} \right) & \text{No decision} \end{cases} \tag{25}$$

where $\tilde{w}_{j'}^k$ is the normalized weight of particle j' within the chosen hypothesis, and the no-decision command (11) is applied whenever no Bayesian decision can be reached within the maximum cost-evaluation horizon $h_{\max}$. The information-enhancement trajectory shaping law (IETS) shares the first four modes of (25), but in the no-decision branch it replaces the saturated linear command with the trajectory-shaping minimization (24), thereby exploiting the decision ambiguity to enhance the estimator's information rather than merely preventing chattering. Explicitly, IETS is given by

$$u_M^k = \begin{cases} +a_M^{\max} & H_1 \text{ decided} \\ \displaystyle\sum_{j' \in H_2} \tilde{w}_{j'}^k \frac{z_1^{j'}(t_{\mathrm{go}}^{j'})}{z_1^*(t_{\mathrm{go}}^{j'})} a_M^{\max} & H_2 \text{ decided} \\ \displaystyle\sum_{j' \in H_3} \tilde{w}_{j'}^k \frac{z_1^{j'}(t_{\mathrm{go}}^{j'})}{z_1^*(t_{\mathrm{go}}^{j'})} a_M^{\max} & H_3 \text{ decided} \\ -a_M^{\max} & H_4 \text{ decided} \\ \arg \displaystyle\min_{u_M^k \in \mathscr{U}_M^k} \det\left((\bar{J}^{k+\bar{h}})^{-1} \right) & \text{No decision.} \end{cases} \tag{26}$$

As noted at the outset of this section, the target parameter θ_T is treated as a known constant and is therefore suppressed from the singular-boundary notation $z_1^*(t_{\mathrm{go}})$; its explicit dependence is reinstated in the robust setting of the next section, where θ_T is no longer known.

4 Robust Estimation-Aware Guidance

The framework reviewed in Section 3 relies critically on knowledge of the target's parameters θ_T. The singular boundary z_1^*, the costs C_{ij} in (15), and the per-hypothesis acceleration commands in (25) all depend on θ_T. When θ_T is uncertain, this dependence introduces a coupling between the parameter posterior, the state posterior, and the optimal decision.

This section presents a robust estimation-aware guidance law (REAGL), which extends the framework to handle the parametric uncertainty sets (2)–(3). The construction proceeds in four steps: 1) augmenting the IMMPF mode set to span the parameter uncertainty, 2) characterizing the resulting family of singular boundaries, 3) redefining the hypotheses on a particle-by-particle basis, and 4) applying the Bayesian decision rule to the joint state-parameter posterior. The section concludes by presenting a robust version of the information-enhancement trajectory shaping law, which extends IETS to the parametric-uncertainty setting.

4.1 Augmented IMMPF Mode Set

Following [10], the target's acceleration command is modeled as a Markov-switching process between two extreme modes, $u_T \in \{+a_T^{\max}, -a_T^{\max}\}$. To capture the additional uncertainty in $a_T^{\max}$, the

IMMPF mode set is augmented with ℓ discrete representative values $\{a_{T,1}^{\max},\ldots,a_{T,\ell}^{\max}\}$ uniformly spanning the interval $[a_{T,1}^{\max},a_{T,2}^{\max}]$. The resulting 2ℓ modes are

$$u_T = \begin{cases} +a_{T,1}^{\max} & n=1 \\ \quad\vdots & \\ +a_{T,2}^{\max} & n=\ell \\ -a_{T,1}^{\max} & n=\ell+1 \\ \quad\vdots & \\ -a_{T,2}^{\max} & n=2\ell. \end{cases} \tag{27}$$

We assume that the target's maximal acceleration capability is constant throughout the engagement. Thus, the TPM $\Pi \in \mathbb{R}^{2\ell\times 2\ell}$ can have at most two nonzero entries per row, corresponding to staying at the same value or transitioning to an acceleration command having the same magnitude but a different sign.

Remark 1. *An extension allowing time-varying acceleration bounds, as assumed in [14] for deterministic cases, is possible by populating additional entries of* Π.

The remaining parameter, τ_T, is more difficult to handle because it is unobservable until the target switches its acceleration command. We therefore adopt a worst-case approach: the interceptor uses the smallest admissible value $\tau_T = \tau_{T,1}$, corresponding to the most agile target in the uncertainty set. This choice yields the smallest singular region in the family parameterized by (2), and hence the most conservative capture-zone characterization. From the interceptor's perspective, this under-promises rather than over-promises its guaranteed-capture set: the assumed singular region is contained in the singular region of the true target, so a particle classified as residing in the singular region under the worst-case τ_T remains in that region also under the actual parameter. The opposite choice, $\tau_T = \tau_{T,2}$, would risk classifying a particle as residing inside the singular region when, under the true (smaller) τ_T, it is in fact outside, leading to commands that fail to prevent the miss.

4.2 Singular Region Under Parameter Uncertainty

For each particle j with associated parameter mode $\theta_T^j = [\tau_{T,1}, a_{T,n^j}^{\max}]^T$, a distinct singular region boundary $z_1^*(t_{\mathsf{go}};\theta_T^j)$ is induced. Figure 4 shows the family of upper boundaries for a representative scenario in which $\tau_T \in \{0.16, 0.20, 0.24\}$ s and $a_T^{\max} \in \{16, 20, 24\}$ g. A more agile target, that is, one having a smaller τ_T or a larger $a_T^{\max}$, yields a smaller singular region, reflecting the fact that a more agile target is harder to capture. The differences between the boundaries grow with t_{go}, which is why the parameter uncertainty matters most early in the engagement.

4.3 Robust Bayesian Decision

In a direct certainty-equivalent strategy, the IMMPF point (MMSE or MAP) estimate of $a_T^{\max}$ would be used in z_1^*, followed by applying the nominal EAGL law of Sec. 3. This approach is suboptimal, as it uses only the first moment of the parameter posterior, which is available from the IMMPF, and not the entire posterior. To circumvent this limitation, we define the four hypotheses on a particle-by-particle basis. Let z_1^j and t_{go}^j denote the ZEM and time-to-go associated with particle j, and let $z_1^*(t_{\mathsf{go}}^j;\theta_T^j)$ be the singular boundary for that particle's parameter mode. Particle j is then assigned to a hypothesis according to the following rule:

$$\begin{aligned} H_1:&\quad z_1^j > z_1^*(t_{\mathsf{go}}^j;\theta_T^j), \\ H_2:&\quad \left|z_1^j\right| \le z_1^*(t_{\mathsf{go}}^j;\theta_T^j) \text{ and } n^j \in \{1,\ldots,\ell\}, \\ H_3:&\quad \left|z_1^j\right| \le z_1^*(t_{\mathsf{go}}^j;\theta_T^j) \text{ and } n^j \in \{\ell+1,\ldots,2\ell\}, \\ H_4:&\quad z_1^j < -z_1^*(t_{\mathsf{go}}^j;\theta_T^j). \end{aligned} \tag{28}$$

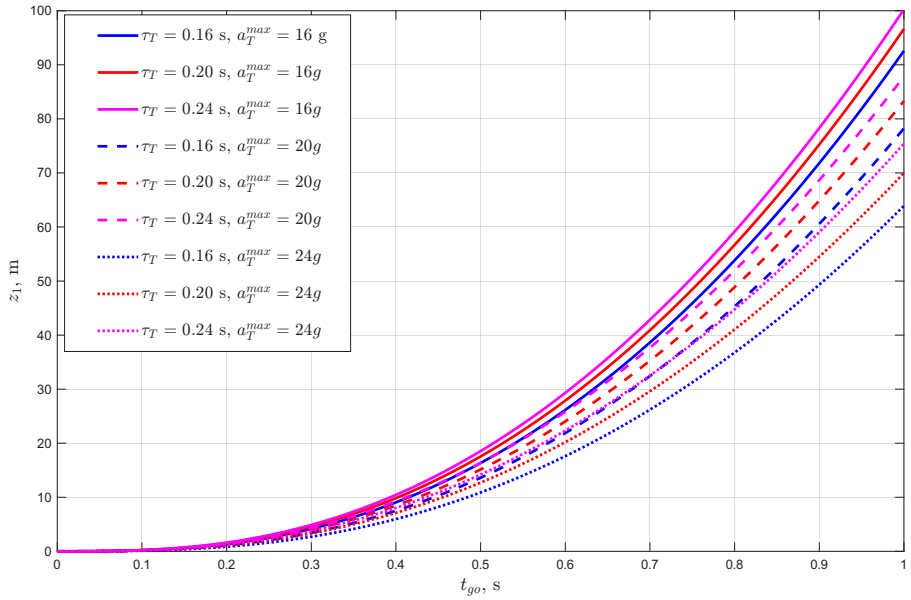


Figure 4 – Upper boundaries of the singular region for an interceptor with $\tau_M = 0.2$ s and $a_M^{\max} = 45$ g, against targets with $\tau_T \in \{0.16, 0.20, 0.24\}$ s and $a_T^{\max} \in \{16, 20, 24\}$ g.

Figure 5 shows the resulting partition for a particle cloud overlaid on three distinct singular boundaries corresponding to three discrete values of $a_T^{\max}$. A larger $a_T^{\max}$ yields a smaller singular region, so the boundaries contract toward the origin as $a_T^{\max}$ increases. Note that particles with identical (z_1, t_{go}) may be assigned to different hypotheses depending on their assumed $a_T^{\max}$, thus capturing the joint posterior over state and parameters.

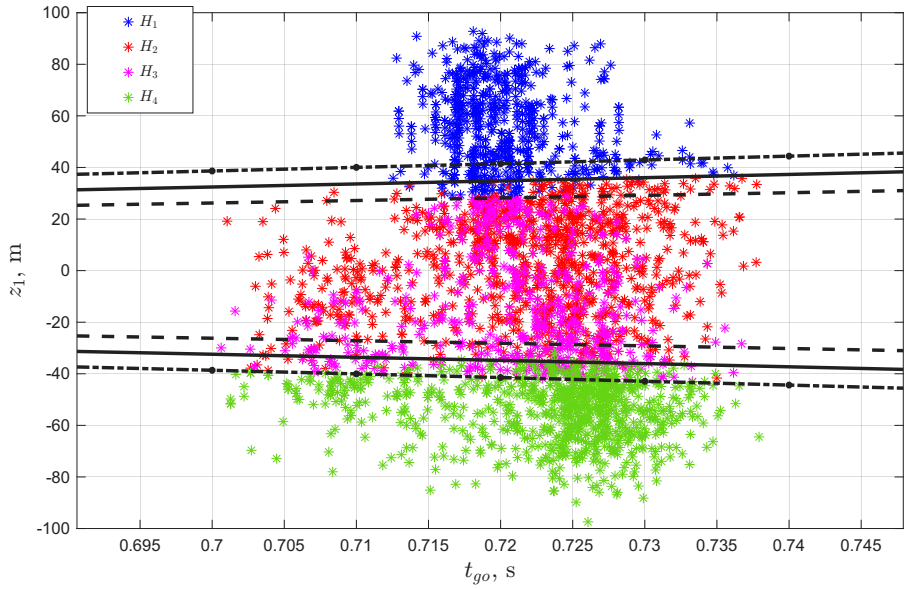


Figure 5 – IMMPF particle cloud, partitioned according to the robust hypotheses (28). Different colors denote particles in H_1, H_2, H_3, H_4. The black dashed, solid, and dashed-dotted lines are the singular region boundaries induced by the discrete values $a_T^{\max} \in \{16, 20, 24\}$ g, respectively, all evaluated at the worst-case time constant $\tau_T = 0.16$ s.

With this hypothesis assignment, the Bayesian decision criterion (15)–(16) applies unchanged. The cost formulas (19) and (20) extend to the robust case by replacing the parameter-independent singular boundary $z_1^*(t_{\text{go}}^{j'})$ with its particle-specific counterpart $z_1^*(t_{\text{go}}^{j'}; \theta_T^{j'})$. Each particle is thus propagated forward using its own dynamic model and contributes to the cost accumulation through the weighted distance to its own singular boundary. The computational reduction conditions of Lemma 1 of [10] ex-

tend directly to the robust setting, with the regions $\mathscr{A}_i$ defined per particle using its parameter mode. Algorithm 1 therefore applies verbatim, with the understanding that all particle-level computations now reference θ_T^j.

4.4 The REAGL Law

The robust GST-compliant guidance law (REAGL) has the same five-mode structure as (25), except that the linear commands for hypotheses H_2 and H_3 now use each particle's own singular boundary, and the no-decision command is the saturated linear command evaluated at the worst-case parameter,

$$u_M^k = \begin{cases} +a_M^{\max} & H_1 \text{ decided} \\ \displaystyle\sum_{j' \in H_2} \tilde{w}_{j'}^k \frac{z_1^{j'}(t_{\mathrm{go}}^{j'})}{z_1^*(t_{\mathrm{go}}^{j'};\, \theta_T^{j'})}\, a_M^{\max} & H_2 \text{ decided} \\ \displaystyle\sum_{j' \in H_3} \tilde{w}_{j'}^k \frac{z_1^{j'}(t_{\mathrm{go}}^{j'})}{z_1^*(t_{\mathrm{go}}^{j'};\, \theta_T^{j'})}\, a_M^{\max} & H_3 \text{ decided} \\ -a_M^{\max} & H_4 \text{ decided} \\ a_M^{\max}\, \mathrm{sat}\left(\dfrac{z_1(t_{\mathrm{go}})}{k z_1^*(t_{\mathrm{go}};\, \theta_T^{\mathrm{wc}})} \right) & \text{No decision} \end{cases}, \tag{29}$$

where $\theta_T^{\mathrm{wc}} \triangleq [\tau_{T,1}, \hat{a}_T^{\max}]^T$ collects the worst-case time constant of Section 4.1 and the IMMPF point estimate $\hat{a}_T^{\max}$ of the acceleration bound. We emphasize, for added clarity, that REAGL departs from EAGL only in the per-particle parameter dependence of the H_2 and H_3 commands and of the singular boundary.

4.5 The RIETS Law

REAGL applies the saturated linear command whenever the Bayesian decision is nonunique. The robust information-enhancement trajectory shaping (RIETS) law replaces that fallback with an information-maximizing command, in direct analogy to the relation between IETS and EAGL in [10]. When the no-decision branch is invoked, the interceptor command is admissible over the set $\mathscr{U}_M^k$ defined in (23). As in [10], this freedom is exploited to maximize the information available to the estimator about the game's state. The no-decision branch of RIETS is therefore

$$u_M^k = \arg \min_{u_M^k \in \mathscr{U}_M^k} \det\left((\bar{J}^{k+\bar{h}})^{-1} \right), \tag{30}$$

where $(\bar{J}^{k+\bar{h}})^{-1}$ is the (z_1, t_{go}) subblock of the inverse FIM, that is, the predicted posterior CRLB, evaluated at the trajectory-shaping horizon $\bar{h}$. The FIM is the same particle-based approximation of [19] used by IETS,

$$J^k \approx -\sum_{j'=1}^{N_p} w_{j'}^k\, \nabla_{\mathbf{x}^k}^2 \log p\left(\mathbf{y}^k \mid \mathbf{x}_{j'}^k \right) \tag{31}$$

now summed over the N_p particles of the augmented IMMPF.

RIETS, which is the robust counterpart of IETS, employs the same information-enhancement objective (30) as IETS; the difference lies not in the optimization but in the objects on which it operates. First, the FIM (31) is formed over the augmented IMMPF's particle cloud (representing the posterior joint density of state and parameters) rather than the particle cloud with perfectly known target parameters of [10], so the predicted CRLB $(\bar{J}^{k+\bar{h}})^{-1}$ reflects the dispersion in particle dynamics induced by the uncertain $a_T^{\max}$: particles carrying different acceleration bounds propagate along different trajectories under a common candidate command, and the minimization is performed against this parameter-spanning ensemble, as shown in Fig. 5. Second, the state resulting from each candidate command is classified against the particle-specific singular boundaries $z_1^*(t_{\mathrm{go}}^j;\, \theta_T^j)$ of (28), rather than the single, parameter-known boundary of [10]. Because the boundary is parameter-dependent, sharpening the estimates of z_1 and t_{go} (the coordinates retained in (30)) sharpens the boundary itself, thereby improving the reliability of the robust hypothesis assignment that drives the decision, which is more

fundamental than merely reducing control chattering. For added clarity, we emphasize that the shaping acts on the state estimate; the target parameters enter only through the augmented ensemble and the per-particle boundaries, and are not themselves the object of the minimization (30).
RIETS, therefore, coincides with REAGL in its first four modes, given in (29), and differs only in the no-decision branch, invoked whenever no unique Bayesian decision can be reached within the maximum cost-evaluation horizon $h_{\max}$: where REAGL applies the saturated linear command, RIETS applies the information-maximizing command (30).

5 Simulation Study

The proposed REAGL and RIETS laws are evaluated in a tactical ballistic missile defense scenario, based on that of [10, 12], except for the introduction of parametric uncertainty in the target's dynamic parameters.

5.1 Simulation Scenario

Both players have constant speeds $V_M = V_T = 2500$ m/s. The interceptor's parameters are $a_M^{\max} = 45$ g and $\tau_M = 0.2$ s. The nominal engagement time is $t_f = 3$ s. Line-of-sight angular measurements are sampled at $f = 100$ Hz with noise standard deviation $\sigma = 0.5$ mrad. Provided by an external radar, the initial state is modeled as Gaussian distributed, $\mathcal{N}(\bar{x}_R, P_R)$, where $\bar{x}_R$ is the true initial radar state, and

$$P_R = \mathrm{diag}\left\{(10^{-3})^2,\ \left(\tfrac{\pi}{180}\right)^2,\ \left(3\tfrac{\pi}{180}\right)^2,\ 100^2\right\} \tag{32}$$

where the $[\rho, \lambda, \gamma_T, V_T]$ ordering of the radar state of [10] is used. The IMMPF is initialized by drawing weighted samples from this Gaussian and propagating them through the nonlinear coordinate transformation to (9).
The target performs a single bang-bang maneuver, that is, a one-time reversal of its acceleration command, with the switch time uniformly distributed over $[1.5, 3]$ s. This interval was found in [12, 10] to be the most challenging time window for the interceptor. The target's parameters are uniformly distributed over the uncertainty set

$$\tau_T \in [0.16, 0.24] \text{ s}, \qquad a_T^{\max} \in [16, 24] \text{ g}, \tag{33}$$

that is, $\pm 20\%$ of their nominal values $\tau_T = 0.2$ s and $a_T^{\max} = 20$ g. We adopt $\pm 20\%$ as a representative uncertainty level, reflecting the operational case in which the target is known only by family and not by its specific airframe.

5.2 Design Considerations

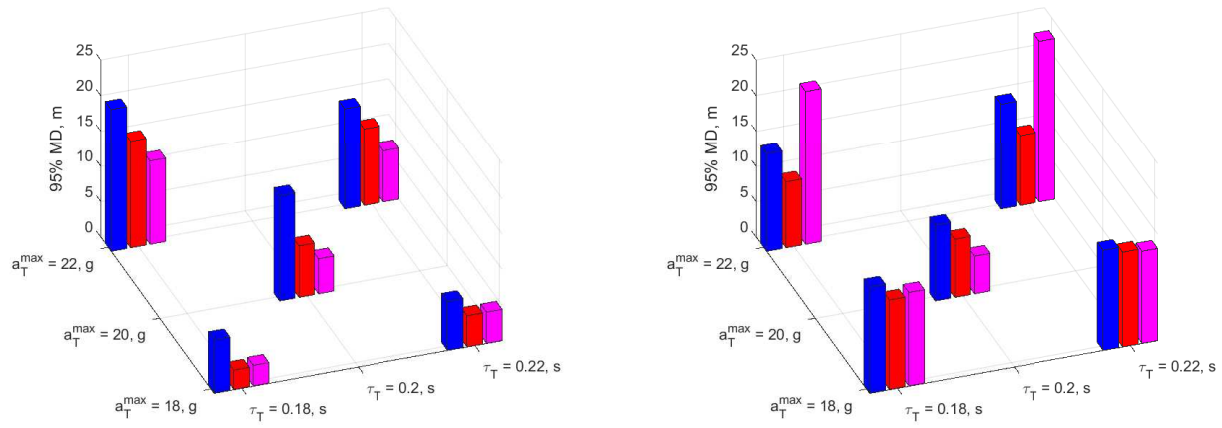
The IMMPF uses $\ell = 3$ discrete acceleration modes, for a total of $2\ell = 6$ modes, with values $\{16, 20, 24\}$ g uniformly spanning the uncertainty interval. Each mode uses 500 particles, for a total particle population of $N_p = 3000$. The worst-case time constant $\tau_T = 0.16$ s is used, as explained in Sec. 4. For all four laws evaluated against the uncertain target, the IMMPF transition probability matrix is taken to be time-invariant, under the simplistic assumption of a single bang-bang maneuver, with $\pi_{12}^k = 2 \times 10^{-2}$ and $\pi_{21}^k = 10^{-3}$.
For the best achievable performance, we use an upper-bound (UB) baseline computation, in which the target's parameters are perfectly known to the interceptor, the IMMPF uses two modes ($\ell = 1$), and the time-varying TPM of [10] is applied,

$$\pi_{12}^k = \begin{cases} 10^{-3} & t_k \le t_s \\ c_{12}\,\dfrac{\beta}{2\alpha\Gamma(1/\beta)}\exp\left(-\dfrac{|t_k - \mu|}{\alpha}\right)^{\beta} & t_k > t_s \end{cases} \tag{34}$$

with $t_s = 1.9$ s, $c_{12} = 0.16$, $\mu = 2.5$, $\beta = 6$, $\alpha = 0.45$, while $\pi_{21}^k = 10^{-3}$ is held constant.
The Bayesian decision criterion uses a maximum cost-evaluation horizon of $h_{\max} = 16/f = 0.16$ s. The trajectory-shaping horizon is set to $\bar{h} = 100$ samples (corresponding to 1 s), and trajectory shaping is disabled in the final 1 s of the engagement, when its lookahead would exceed the remaining flight time. The threshold for the admissible set (23) is $W_{\text{Thres}} = 0.15$. The DGL1 saturated linear command in $\mathscr{D}_0$ uses gain $k = 0.7$, as in [10].

5.3 Sensitivity to Parameter Mismatch

To illustrate the limitation that motivates the robust approach, we examine how the nominal estimation-aware laws respond to a mismatch between the assumed and the true target parameters. The regular DGL1, EAGL, and IETS laws are evaluated against a target whose parameters are fixed at off-nominal values on a grid spanning $\pm 10\%$ about the nominal parameters. Figure 6 reports the 95th percentile miss distance for two switch times of the bang-bang maneuver: 1) a mid-engagement maneuver reversal at $t_{sw} = 2$ s, and 2) a maneuver at the end of the engagement, $t_{sw} = 3$ s, which means that the target performs a constant-acceleration maneuver (no maneuver reversal) throughout the interception. When the target reverses maneuver in mid-engagement, Fig. 6a, IETS performs best across the grid, retaining the advantage it holds in [10] even under parameter mismatch. However, against a constant maneuver by a higher-than-nominal capability target, Fig. 6b, its miss distance grows sharply, whereas EAGL preserves its improvement over the regular DGL1. These two outcomes share a single underlying mechanism. IETS drives the game state to reside on the boundary of the singular region, a boundary evaluated assuming known parameters. When the true acceleration bound is larger, the singular region is smaller than assumed. Held against the nominal boundary, the state then falls outside the true region, rendering the resulting guidance decisions significantly erroneous. A mid-engagement reversal then frees the state to return inside the singular region, and the shaping recovers its advantage. In contradistinction, under a constant maneuver the state remains outside the singular region for the rest of the engagement, producing the large miss distance of Fig. 6b. We conclude that IETS is not a weak law; rather, it is a non-robust one: it performs best precisely when the target maneuvers as the shaping anticipates, and worst when a higher-capability target denies it that opportunity. For this reason the nominal IETS law is not carried into the ensuing guidance laws comparison of Sec. 5.4. To robustly address IETS's sensitivity to parameter mismatch, the RIETS law of Sec. 4.5 shapes against the particle-specific boundaries $z_1^*(t_{\text{go}}^j; \theta_T^j)$ of the augmented ensemble rather than against a single, nominal boundary, thus seeking the benefit of trajectory shaping while avoiding its fragility due to parameter mismatch.



(a) $t_{sw} = 2$ s. (b) $t_{sw} = 3$ s.

Figure 6 – Sensitivity to parameter mismatch: 95th percentile miss distance over a grid of off-nominal target parameters, for a maneuver reversal at (a) $t_{sw} = 2$ s and at (b) $t_{sw} = 3$ s (effectively no maneuver reversal). Comparison of the regular DGL1 (blue), EAGL (red), and IETS (magenta) guidance laws.

5.4 Monte Carlo Performance

A 4000-run Monte Carlo simulation is conducted comparing four guidance laws under the parametric uncertainty defined in (33). For a fair comparison, all four laws share the same augmented IMMPF (the $2\ell = 6$ mode configuration of Section 5.2); they differ only in how its posterior is used to select the command:

1. Classical DGL1 [9], applied in a certainty-equivalent fashion: the singular boundary z_1^* is evaluated at the MMSE estimate $\hat{a}_T^{\max}$, and the command follows from the position of the MMSE state estimate relative to that boundary.

2. EAGL [10], which applies its Bayesian decision criterion to the state posterior but evaluates the singular boundary z_1^* at the MMSE estimate $\hat{a}_T^{\max}$, in place of the true, unknown target parameter. This implementation uses only the first moment of the parameter posterior, provided by the augmented IMMPF, and not the entire posterior.

3. REAGL, as presented in Sec. 4.4.

4. RIETS, as presented in Sec. 4.5.

For reference, three upper bounds are also computed assuming perfect knowledge of the target's parameters: the DGL1 UB; the EAGL/REAGL UB, the two coinciding by construction, as discussed in Section 4.4; and the IETS/RIETS UB, which coincide for the same reason.
Figure 7 shows the empirical miss distance cumulative distribution functions (CDFs). Table 1 summarizes the warhead lethality radii required to achieve a single-shot kill probability (SSKP) of 0.9, denoted R_{90}, and the SSKPs achieved by a warhead of a fixed lethality radius of 10 m.

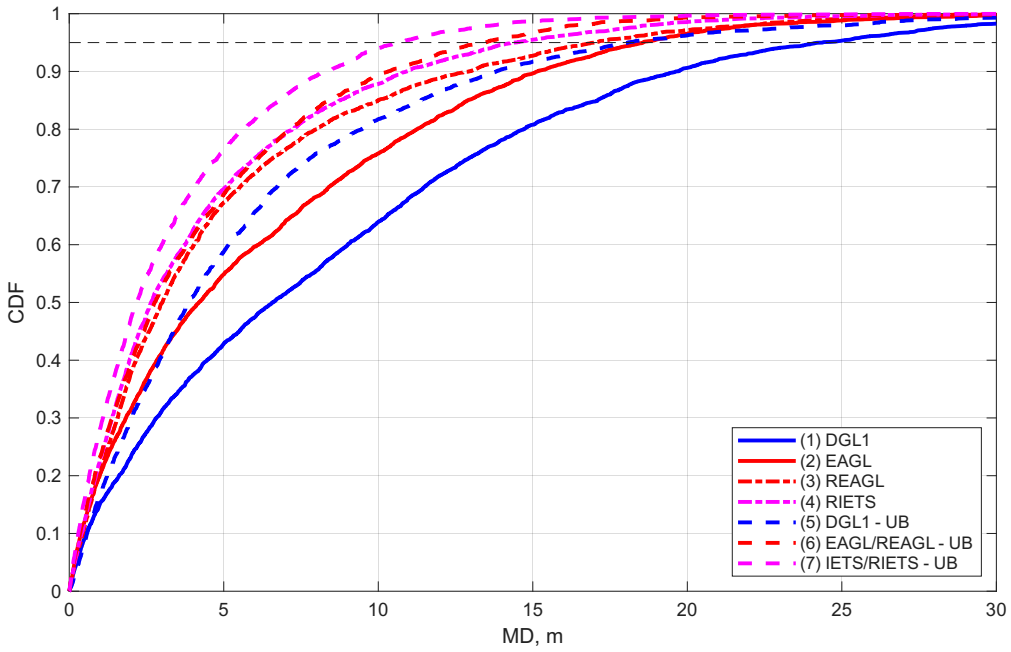


Figure 7 – Miss distance CDF comparison (4000 Monte Carlo runs). Lines (1)–(4) are evaluated under parametric uncertainty: solid lines, (1) DGL1 and (2) EAGL; dashed-dotted lines, (3) REAGL and (4) RIETS. Dashed lines are the upper bounds, computed with the true target parameters: (5) DGL1, (6) EAGL/REAGL, (7) IETS/RIETS.

The results separate the gain due to the estimation-aware decision logic from the gain due to robustness to parametric uncertainty. The three upper bounds in Table 1 and Fig. 7, computed with the target parameters known, are reference ceilings; they are not on equal footing with the curves obtained under uncertainty, but among themselves they form a like-for-like comparison. The DGL1 upper bound (Line 5) requires $R_{90} = 13.81$ m, whereas the EAGL/REAGL upper bound (Line 6), driven by the same IMMPF, requires only 10.32 m. Since the two laws share the same estimator, this reduction is attributable solely to the estimation-aware Bayesian decision criterion, which acts on the entire posterior provided by the IMMPF rather than on its first moment, as the certainty-equivalent DGL1 does. Adding trajectory shaping lowers the IETS/RIETS upper bound (Line 7) further, to 8.23 m.
Under parametric uncertainty, the four laws form a cumulative sequence in which each adds one ingredient to its predecessor; all are evaluated against the same augmented IMMPF and uncertainty set, and are therefore directly comparable. The certainty-equivalent DGL1 (Line 1) degrades severely,

Table 1 – Performance comparison: warhead lethality radius required to guarantee SSKP of 0.9 (R_{90}), and SSKP for a lethality radius of 10 m. UB denotes the upper bound computed with perfect knowledge of the target parameters.

Guidance Law	Lethality Radius (R_{90})		SSKP (10 m Lethality Radius)	
	Value (m)	UB (m)	Value	UB
DGL1	19.64	13.81	0.639	0.817
EAGL	15.15	10.32	0.758	0.893
REAGL	12.95	10.32	0.850	0.893
RIETS	10.90	8.23	0.878	0.940

requiring $R_{90} = 19.64$ m. Applying the estimation-aware decision criterion, while still using the point estimate of $a_T^{\max}$ for the singular boundary, EAGL (Line 2) reduces this to 15.15 m. Replacing the point estimate with the joint state-parameter posterior, REAGL (Line 3) reaches 12.95 m, a 34.1% improvement over DGL1. Finally, adding active trajectory shaping in the no-decision branch, RIETS (Line 4) outperforms its competing guidance laws: it requires $R_{90} = 10.90$ m, the lowest required radius, while providing an SSKP of 0.878 at the 10 m lethality radius, the highest SSKP.
Notably, RIETS under parameter uncertainty comes within 6% in R_{90}, and within 0.015 in SSKP, of the EAGL/REAGL upper bound computed with perfect parameter knowledge. This compares an operational law against an idealized reference, rather than two laws on equal footing; read in that light, this indicates that the robust decision and active shaping together nearly offset, in miss distance terms, the performance ordinarily lost to parametric uncertainty.

6 Conclusions

We have presented an extension of the unified estimation-guidance framework of [10] to the case of uncertain target dynamic parameters. This is a setting of practical importance in tactical ballistic missile defense based on advanced, differential game based guidance laws. In these laws, the singular region boundary, which determines the guidance command, depends on parameters that cannot be assumed known.
The robust estimation-aware guidance law (REAGL) replaces the certainty-equivalent treatment of target parameters with a fully Bayesian decision criterion that operates on the joint state-parameter posterior produced by an augmented IMMPF. The robust information-enhancement trajectory shaping law (RIETS) exploits the nonuniqueness of the optimal Bayesian decision to sharpen the state estimate in the directions most relevant to the parameter-dependent guidance command.
A 4000-run Monte Carlo simulation in a representative ballistic missile defense scenario, with $\pm 20\%$ uniform uncertainty in both τ_T and $a_T^{\max}$, demonstrates that REAGL improves the required warhead lethality radius by 34.1% over the classical DGL1, and that adding trajectory shaping via RIETS yields a further improvement. Adding the information-enhancement branch to REAGL, the RIETS law attains the lowest required lethality radius and the highest single-shot kill probability of the laws evaluated under parametric uncertainty. These results validate that closing the loop between estimation and guidance is critical when realistic parameter uncertainty is present.

7 Contact Author Email Address

For correspondence regarding this paper, please email liraz109@gmail.com.

8 Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] H. S. Witsenhausen. Separation of estimation and control for discrete time systems. *Proceedings of the IEEE*, 59(11):1557–1566, 1971.

[2] J. Shinar and V. Turetsky. What happens when certainty equivalence is not valid? is there an optimal estimator for terminal guidance? *IFAC Proceedings Volumes*, 36(8):175–184, 2003.

[3] J. L. Speyer. An adaptive terminal guidance scheme based on an exponential cost criterion with application to homing missile guidance. *IEEE Transactions on Automatic Control*, 21(3):371–375, 1976.

[4] G. Hexner and T. Shima. Stochastic optimal control guidance law with bounded acceleration. *IEEE Transactions on Aerospace and Electronic Systems*, 43(1):71–78, 2007.

[5] G. Hexner, T. Shima, and H. Weiss. LQG guidance law with bounded acceleration command. *IEEE Transactions on Aerospace and Electronic Systems*, 44(1):77–86, 2008.

[6] J. Shinar, V. Turetsky, and Y. Oshman. Integrated estimation/guidance design approach for improved homing against randomly maneuvering targets. *Journal of Guidance, Control, and Dynamics*, 30(1):154–161, 2007.

[7] I. G. Shaviv and Y. Oshman. Estimation-guided guidance and its implementation via sequential Monte Carlo computation. *Journal of Guidance, Control, and Dynamics*, 40(2):402–417, 2017.

[8] S. Gutman. On optimal guidance for homing missiles. *Journal of Guidance and Control*, 2(4):296–300, 1979.

[9] J. Shinar. Solution techniques for realistic pursuit-evasion games. In C. T. Leondes, editor, *Control and Dynamic Systems*, volume 17, pages 63–124. Academic Press, New York, 1981.

[10] L. Mudrik and Y. Oshman. Unified estimation–guidance framework based on Bayesian decision theory. *Journal of Guidance, Control, and Dynamics*, 2026. Article in Advance.

[11] H. A. P. Blom and E. A. Bloem. Exact Bayesian and particle filtering of stochastic hybrid systems. *IEEE Transactions on Aerospace and Electronic Systems*, 43(1):55–70, 2007.

[12] V. Shaferman and Y. Oshman. Stochastic cooperative interception using information sharing based on engagement staggering. *Journal of Guidance, Control, and Dynamics*, 39(9):2127–2141, 2016.

[13] V. Shaferman. Near-optimal evasion from pursuers employing modern linear guidance laws. *Journal of Guidance, Control, and Dynamics*, 44(10):1832–1844, 2021.

[14] T. Shima and J. Shinar. Time-varying linear pursuit-evasion game models with bounded controls. *Journal of Guidance, Control, and Dynamics*, 25(3):425–432, 2002.

[15] S. Gutman, O. Goldan, and S. Rubinsky. Guaranteed miss distance in guidance systems with bounded controls and bounded noise. *Journal of Guidance, Control, and Dynamics*, 35(3):816–823, 2012.

[16] H. L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory*. Wiley, Hoboken, NJ, 2004.

[17] H. A. P. Blom and Y. Bar-Shalom. The interacting multiple model algorithm for systems with Markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, 1988.

[18] Y. Oshman and P. Davidson. Optimization of observer trajectories for bearings-only target localization. *IEEE Transactions on Aerospace and Electronic Systems*, 35(3):892–902, 1999.

[19] A. Tulsyan, B. Huang, R. B. Gopaluni, and J. F. Forbes. A particle filter approach to approximate posterior Cramér–Rao lower bound: The case of hidden states. *IEEE Transactions on Aerospace and Electronic Systems*, 49(4):2478–2495, 2013.

[20] P. Tichavský, C. H. Muravchik, and A. Nehorai. Posterior Cramér–Rao bounds for discrete-time nonlinear filtering. *IEEE Transactions on Signal Processing*, 46(5):1386–1396, 1998.