

Article

Dynamic Time-Window Nash Equilibrium Strategies for Spacecraft Pursuit–Evasion Games Under Incomplete Strategies

Lei Sun ^{1,*} , Zengliang Han ¹ , Yuhui Wang ¹ , Binpeng Tian ^{2,3} and Panxing Huang ^{4,5}

¹ College of Automation Engineering, Nanjing University of Aeronautics and Astronautics, Nanjing 211106, China; hanzengliang@nuaa.edu.cn (Z.H.); wangyh@nuaa.edu.cn (Y.W.)

² School of Aeronautics, Northwestern Polytechnical University, Xi'an 710072, China; 13031628966@163.com

³ The First Aircraft Institute, Aviation Industry Corporation of China, Ltd., Xi'an 710089, China

⁴ Beijing Institute of Control Engineering, Beijing 100094, China; huangpanxing2006@126.com

⁵ National Key Laboratory of Space Intelligent Control, Beijing 100094, China

* Correspondence: sl0826@nuaa.edu.cn

Abstract

Spacecraft pursuit–evasion in contested environments is complicated by strategic incompleteness: the evader can switch maneuvering modes and deploy multi-domain countermeasures that degrade the pursuer's perception, leading to non-stationary information and distributionally ambiguous interference statistics. A dynamic time-window Nash equilibrium framework is developed for linearized Local Vertical Local Horizontal (LVLH) relative motion under interference-induced uncertainty. Perceptual degradation is modeled via an evidence–theoretic belief representation, and the Jensen–Shannon (JS) divergence is introduced to quantify discrepancies between nominal and interference-corrupted beliefs. The divergence metric drives an adaptive time-window partitioning policy and an uncertainty-aware running cost that balances nominal performance objectives with robustness regularization during high-degradation intervals. In each time window, sufficient conditions are provided for the existence of a local Nash equilibrium, and equilibrium strategies are characterized by the Hamilton–Jacobi–Bellman–Isaacs (HJBI) equation. A global consistency result is established: assuming state continuity, additive cost decomposition, and dynamic-programming compatibility at window boundaries, concatenating the window-wise equilibria yields a Nash equilibrium over the entire horizon. Unlike conventional receding-horizon differential games with a fixed replanning grid, the proposed policy partitions the horizon online in response to perceptual-degradation events and stitches adjacent windows through a continuation value. This boundary stitching enables the global consistency guarantee under additive costs and state continuity. To hedge against ambiguity in interference intensity, a variational distributionally robust optimization (DRO) problem with moment-constrained ambiguity sets is formulated, and the dual worst-case distribution is derived. The resulting Karush–Kuhn–Tucker (KKT) system is reformulated as a finite-dimensional variational inequality, for which an accelerated Alternating Direction Method of Multipliers (ADMM) operator-splitting solver is proposed for efficient real-time computation. Numerical simulations validate the framework and demonstrate improved robustness and computational scalability under time-varying interference compared with fixed-window baselines.



Academic Editor: Carmine Maria Pappalardo

Received: 28 January 2026

Revised: 18 February 2026

Accepted: 23 February 2026

Published: 2 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: spacecraft pursuit–evasion; differential games; Nash equilibrium; dynamic time-windows; incomplete strategies; distributionally robust optimization

1. Introduction

Spacecraft pursuit–evasion (P–E) engagements constitute a central problem class in adversarial proximity operations, with direct implications for interception, inspection, defense of high-value assets, and protection of orbital infrastructure in increasingly contested environments [1]. Unlike cooperative rendezvous and formation-flying missions, P–E interactions are shaped by explicit strategic opposition: the pursuer is driven to reduce separation and achieve capture under fuel, thrust, and safety constraints, whereas the evader is driven to preserve separation, delay interception, or induce mission failure [2]. Relative motion in near-circular orbits is commonly described in the Local Vertical Local Horizontal (LVLH) frame using linearized dynamics, while extensions to non-circular reference orbits and energy-optimal relative-motion formulations provide complementary modeling and optimization foundations for realistic engagements [3].

Typical deployment scenarios include adversarial proximity operations such as on-orbit inspection, rendezvous-and-intercept, and defense of high-value assets, where the pursuer must approach under stringent fuel and safety constraints while the evader actively degrades tracking quality [4]. In such settings, strategic incompleteness arises not only from maneuver-mode switching, but also from deliberate sensing denial (e.g., jamming/spoofing, infrared deception, electro-optical degradation), which corrupts the pursuer’s inference of the evader’s intent and renders the information pattern non-stationary over the engagement.

In contested pursuer evader engagements, the main challenge is not only control under random disturbances. It is also a strategic interaction in which the evader may intentionally influence what the pursuer can observe and believe [5]. In practice, the evader can use countermeasures across multiple domains, such as electromagnetic interference, infrared deception, and degraded electro optical sensing. These actions can corrupt measurements, reduce tracking quality, and cause the pursuer to misread the evader’s maneuver intent. Related anti-deception and anti-jamming methods have also been studied in sensing systems for adversarial settings [6]. At the same time, the evader may switch among different maneuver modes as the relative geometry changes and as it reacts to the pursuer’s behavior. Examples include energy limited drifting, short high-effort evasive bursts, and linear quadratic guidance laws. These factors are tightly linked. Interference reduces state estimation accuracy and also weakens the pursuer’s ability to identify the evader’s current strategy mode. If the mode inference is wrong, the pursuer’s prediction, planning horizon, and control allocation can all change. From an operational view, the engagement is further affected by resource management decisions, such as when to use countermeasures and how to allocate sensing resources. This naturally motivates decision theoretic formulations under partial observability [7]. Finally, robust relative pose and state estimation remains essential for proximity operations when perception is degraded. Therefore, practical estimation pipelines should be considered together with game theoretic guidance methods [8].

The significance of this research is both practical and theoretical. From a practical perspective, guidance policies designed for normal sensing conditions can become unreliable when perception degrades. This risk is especially high in short capture windows, where a delayed response can permanently remove the chance to intercept. From a theoretical perspective, degraded information weakens assumptions that are often used to simplify differential game decision making. These include using stationary strategy classes, evaluating payoffs as time invariant, and assuming known probabilistic disturbance models. When these issues are not addressed, two common problems arise: (i) the decision loop adapts too slowly to sudden drops in information quality, or (ii) it adapts too aggressively with short-sighted control updates that break time consistency and reduce overall perfor-

mance [9]. These points motivate pursuer evader methods that connect online uncertainty quantification with game-consistent decision making, while remaining computationally feasible. Related analytical work in orbital pursuer evader settings also shows that outcomes can be highly sensitive to orbit geometry and available maneuver authority. This further supports the need for frameworks that remain dependable across different operating regimes [10]. Finally, safety and reliability requirements, including constraint satisfaction under uncertainty, motivate architectures that combine planning and control with robustness mechanisms [11].

From a practical standpoint, directly computing equilibrium policies over long horizons is often infeasible in real time, even before accounting for interference-induced uncertainty [12]. When incomplete strategies and non-stationary information are introduced, the computational burden increases further because the pursuer must reason over multiple plausible evader behaviors and potentially unknown interference statistics. Consequently, receding-horizon approximations are frequently adopted to reduce computation and incorporate newly acquired measurements; fast real-time implementations and robust data-driven variants have been actively investigated in the control literature [13]. Beyond deterministic rolling-horizon control, robust sequential decision formulations (e.g., robust Markov decision processes) offer a principled way to encode model ambiguity and adversarial uncertainty [14], and sequential measurement models further clarify how information accumulation affects optimal decision making [15]. In addition, dual-control-style approaches that explicitly trade off task performance and active uncertainty reduction provide another perspective on coupling information gathering with safe interaction planning [16]. Nevertheless, fixed rolling windows introduce their own vulnerabilities: the window length becomes a tuning knob that strongly affects responsiveness, stability, and solution quality, and the resulting sequence of local solutions may not compose into a globally meaningful equilibrium.

Existing research related to P-E and robust decision making has achieved important progress across several directions [17]. Robust control techniques address bounded disturbances and modeling errors, while stochastic and risk-sensitive methods incorporate probabilistic uncertainty. Despite these advances, three limitations remain prominent in contested-space settings [18].

First, many pursuer evader formulations assume that the evader's strategy set is known and does not change over the planning horizon. In real engagements, the evader can switch maneuver logic, and the pursuer's estimate of the evader's current mode can also be distorted by countermeasures [19]. In this setting, relying on a single model can create consistent prediction errors and lead to poor control decisions. Recent work in the wider pursuit-evasion literature points in the same direction. Examples include nature-inspired pursuit-evasion control, learning-based policies, hierarchical option-based maneuver selection, and online identification or adaptation against unknown pursuers. These developments suggest that it is both feasible and necessary to account for non-stationary behavior and uncertainty in the evader's mode [20].

Second, uncertainty is frequently represented by fixed parametric noise distributions or simplified independent-noise models. Under active interference, uncertainty is neither stationary nor purely aleatoric: its statistics vary across time windows, may be correlated, and can be manipulated adversarially [21]. In such cases, committing to a single assumed distribution yields overconfident decisions and poor robustness. A more appropriate approach is to explicitly acknowledge distributional ambiguity and optimize against a family of plausible distributions; moment-based ambiguity sets, tractable distributionally robust optimization (DRO) approximations, and Wasserstein-metric DRO with performance guarantees provide widely used tools toward this direction [22].

Third, the computational resolution in many rolling-horizon schemes is fixed a priori. Fixed windows can be inefficient: long windows may be unnecessary and expensive under benign conditions, whereas short windows may be insufficiently responsive under severe interference or rapid strategy switching. More critically, naive windowing can compromise time consistency: local solutions computed in different windows may not compose into a globally meaningful equilibrium, making performance dependent on heuristic choices rather than principled guarantees [23].

Motivated by these gaps, a unified framework is developed to integrate (i) online uncertainty quantification under incomplete strategies, (ii) adaptive computational resolution through dynamic time windows, and (iii) distributional robustness to ambiguous interference statistics, while preserving rigorous game-theoretic semantics. The key architectural idea is to partition the engagement horizon into sequential time windows whose lengths are adjusted online according to the pursuer's current information quality and interference intensity. Under benign conditions, longer windows enable more strategic look-ahead and reduce unnecessary recomputation; under severe perception degradation, shorter windows enable rapid policy refresh and tighter coupling to newly accumulated evidence. Importantly, the dynamic time-window mechanism is embedded into a window-wise equilibrium construction in which local Nash equilibria are computed conditioned on boundary information and a continuation value, thereby avoiding purely heuristic replanning.

A central modeling component is the representation of strategic incompleteness. An evidence-theoretic belief model is adopted for evader strategy modes, enabling epistemic uncertainty to be expressed and updated beyond classical probability assignments [24]. A divergence-based metric derived from these beliefs is introduced as a real-time indicator of perceptual degradation; in particular, the Jensen–Shannon divergence provides a bounded and symmetric discrepancy measure that is well-suited for online monitoring and threshold-based triggering [25]. This indicator is used to drive both window adaptation and an uncertainty-aware cost evaluation so that the relative weighting between nominal performance and robustness is adjusted consistently with the observed sensing quality.

Finally, the overall framework is tailored to real-time computation. The resulting optimality conditions admit operator-splitting implementations with inexpensive projections onto simple constraint sets, and an accelerated ADMM-type solver is developed so that window-wise equilibrium computation can be performed within controllable iteration budgets [26]. This computational design aligns naturally with the dynamic time-window architecture: when the window shrinks due to high degradation, the solver budget can be kept modest while still achieving timely policy updates; when the window expands in calmer phases, additional iterations can be allocated to refine the equilibrium solution. More broadly, the trend toward onboard autonomy and AI-enabled spacecraft operations further motivates efficient and interpretable decision pipelines that can be executed under constrained resources and uncertain sensing [27,28].

Different from fixed-window or heuristic rolling-horizon pursuit–evasion schemes that re-solve local games on a preselected horizon, our framework (i) triggers window boundaries by online belief/state signals induced by perceptual degradation, (ii) enforces boundary consistency through an explicit continuation value, and (iii) embeds a window-wise VDRO layer to hedge against distributional ambiguity in interference intensity. As a result, adaptivity, global equilibrium consistency, and distributional robustness are achieved in a unified design rather than being addressed in isolation.

The main contributions and innovations of this work are summarized as follows:

- Dynamic time-window equilibrium architecture under incomplete strategies. An adaptive time-window mechanism is developed to link real-time perceptual degra-

dation and interference intensity to computational resolution, enabling a principled trade-off between responsiveness and planning depth.

- Local equilibrium formulation with global consistency guarantees. Sufficient conditions are established for the existence of local Markovian Nash equilibria within each time window, and a global consistency result is provided to show that the concatenation of window-wise equilibria forms a Nash equilibrium over the full engagement horizon under mild structural conditions.
- Distributional robustness to ambiguous interference statistics. A variational distributionally robust game is formulated using moment-constrained ambiguity sets for interference intensity, yielding a principled hedge against distributional misspecification in adversarial sensing environments.
- VI reformulation and accelerated ADMM solver for real-time computation. The VDRO KKT system is reformulated as a finite-dimensional VI, and an accelerated ADMM operator-splitting method with simple projections and controllable iteration budgets is designed for efficient window-wise equilibrium computation.

The remainder of this paper is organized as follows. Section 2 formulates the spacecraft relative motion dynamics and the pursuit–evasion game model under strategic incompleteness. Section 3 introduces the dynamic time-window mechanism and presents the window-wise Nash equilibrium formulation. Section 4 establishes global equilibrium consistency and develops the variational distributionally robust optimization framework. Section 5 presents the variational inequality reformulation of the KKT conditions and the accelerated ADMM solver. Section 6 reports numerical simulations. Section 7 concludes the paper and outlines future research directions.

2. Problem Formulation and Game Framework

2.1. Relative Motion Dynamics in LVLH Coordinate System

The pursuit–evasion engagement between spacecraft is formulated within the Local Vertical Local Horizontal (LVLH) coordinate system centered at a reference satellite. Following the Clohessy–Wiltshire equations for near-circular orbits, the relative dynamics between the pursuer P and evader E are expressed as:

$$\begin{aligned}\ddot{x} - 2\omega\dot{y} - 3\omega^2x &= u_x \\ \ddot{y} + 2\omega\dot{x} &= u_y \\ \ddot{z} + \omega^2z &= u_z\end{aligned}\quad (1)$$

where $[x, y, z]^T$ represents the relative position vector, ω denotes the orbital angular velocity of the reference orbit, and $[u_x, u_y, u_z]^T$ represents the control acceleration vector.

Defining the state vector $\mathbf{X} = [x, y, z, \dot{x}, \dot{y}, \dot{z}]^T$ and control input $\mathbf{U} = [u_x, u_y, u_z]^T$, the system dynamics can be rewritten in state-space form:

$$\dot{\mathbf{X}} = \mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{U}\quad (2)$$

with system matrices:

$$\mathbf{A} = \begin{bmatrix} \mathbf{0}_{3\times 3} & \mathbf{I}_{3\times 3} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{0}_{3\times 3} \\ \mathbf{I}_{3\times 3} \end{bmatrix}\quad (3)$$

where

$$\mathbf{A}_{21} = \begin{bmatrix} 3\omega^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\omega^2 \end{bmatrix}, \quad \mathbf{A}_{22} = \begin{bmatrix} 0 & 2\omega & 0 \\ -2\omega & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}\quad (4)$$

To clarify the geometry of the Local Vertical Local Horizontal (LVLH) frame and the meaning of the state components used in (1)–(4), Figure 1 depicts the radial (x), along-track (y), and cross-track (z) axes together with the relative position vector between the evader and the pursuer. This coordinate convention is adopted throughout the paper for defining the relative state, the capture condition, and the subsequent window-wise game formulation.




Figure 1. LVLH relative-motion frame. The state uses the relative position components (x, y, z) along the radial, along-track, and cross-track axes, respectively, and the relative distance is $r = \sqrt{x^2 + y^2 + z^2}$. The capture condition is defined by $r \leq r_{\text{capt}}$, which is encoded by the terminal cost in (7).

Remark 1 (Model validity). *The linearized LVLH relative-motion model is adopted under standard proximity-operation assumptions: the reference orbit is near-circular with slowly varying mean motion, and the relative separation and relative velocity remain sufficiently small compared with the orbital radius so that higher-order nonlinear terms can be neglected. Accordingly, the proposed analysis and synthesis are intended for close-range engagements in which the linearization error stays bounded over each planning window.*

2.2. Differential Game Model for Spacecraft Pursuit–Evasion

The pursuit–evasion engagement between two spacecraft is formulated as a two-player, zero-sum differential game, denoted by the tuple:

$$\mathcal{G} = (P, E; \mathbf{X}, f(\mathbf{X}); U_P, U_E; J), \quad (5)$$

where P denotes the pursuer spacecraft, E denotes the evader spacecraft, and $\mathbf{X} \in \mathbb{R}^n$ is the system state vector. The game evolves over a fixed time horizon $t \in [0, T]$. The control strategies $U_P : \mathbb{R}^n \rightarrow \mathbf{U}_P$ and $U_E : \mathbb{R}^n \rightarrow \mathbf{U}_E$ are the mappings from the state space to compact, convex control sets $\mathbf{U}_P \subset \mathbb{R}^{m_P}$ and $\mathbf{U}_E \subset \mathbb{R}^{m_E}$, respectively.

Both players aim to optimize a common cost functional J over the engagement horizon. The pursuer P seeks to minimize J , while the evader E attempts to maximize it:

$$J(\mathbf{X}, U_P, U_E) = \int_0^T G(\mathbf{X}(t), U_P(t), U_E(t)) dt + H(\mathbf{X}(T)), \quad (6)$$

where $G(\cdot)$ is the instantaneous running cost and $H(\cdot)$ is the terminal cost.

The instantaneous cost penalizes both the distance penalty and control energy:

$$G(\mathbf{X}, U_P, U_E) = \underbrace{\alpha r^2}_{\text{Distance Penalty}} + \underbrace{\beta_P \sum_{i=1}^3 U_{P_i}^2 + \beta_E \sum_{j=1}^2 U_{E_j}^2}_{\text{Control Energy}}, \quad (7)$$

where $\alpha > 0$ is a weighting factor for relative distance minimization, and $\beta_P, \beta_E > 0$ penalize the control energy expenditure of the pursuer and evader, respectively.

The exponential terminal form in (7) is chosen to provide a smooth and bounded success–failure discriminator: it is continuously differentiable with respect to $r(T)$, which is convenient for the HJBI-based equilibrium characterization, while still producing a sharp transition around r_{capt} . In Section 6.7, the sensitivity of performance to (η, r_{capt}) is quantified to confirm that the conclusions are not tied to a single tuning.

The terminal cost encodes the capture outcome:

$$H(\mathbf{X}(T)) = \eta \left(1 - 2 \exp \left(-\frac{r^2(T)}{r_{\text{capt}}^2} \right) \right). \quad (8)$$

where, $\eta > 0$ scales the terminal payoff, and r_{capt} is the capture radius defining successful interception. Since $d_{\text{JS}} \in [0, \ln 2]$, it is convenient to use the normalized index $\bar{d}_{\text{JS}} = d_{\text{JS}} / (\ln 2) \in [0, 1]$ to interpret degradation levels. In our implementation, $\bar{d}_{\text{JS}} < 0.2$ is treated as low degradation, $0.2 \leq \bar{d}_{\text{JS}} \leq 0.6$ is treated as moderate degradation, and $\bar{d}_{\text{JS}} > 0.6$ is treated as high degradation, which guides threshold selection (e.g., d_{thresh}) and window shrinking.

2.3. Modeling Strategic Incompleteness

In practical engagement scenarios, the evader employs multi-domain countermeasures (e.g., electromagnetic jamming, infrared decoys) to actively degrade the pursuer’s situational awareness. This results in incomplete information regarding the evader’s true strategy. The strategic incompleteness arises from two key factors:

1. **Multi-mode Strategy Switching:** The evader can dynamically switch among a set of distinct guidance strategies, such as linear-quadratic optimal evasion or maximum-effort maneuvers.
2. **Active Sensor Interference:** The evader deploys countermeasures that corrupt sensor measurements, introducing non-stationary, correlated observation noise that violates conventional estimation assumptions.

To formally characterize the resultant perceptual uncertainty, we adopt an evidence-theoretic framework. Let $\Theta = \{\theta_1, \theta_2, \dots, \theta_M\}$ be the frame of discernment containing M possible evader strategy modes. The uncertainty is quantified via a Basic Probability Assignment (BPA) function $m : 2^\Theta \rightarrow [0, 1]$, which is transformed into a decision-ready probability distribution via the Pignistic transformation $\text{BetP}(\theta_i)$.

The discrepancy between the estimated belief distribution (under interference) and the nominal (interference-free) distribution is measured using the Jensen–Shannon (JS) divergence [29]:

$$d_{\text{JS}}(m_1, m_2) = \frac{1}{2} [D_{\text{KL}}(\text{BetP}_1 \| \mathcal{M}) + D_{\text{KL}}(\text{BetP}_2 \| \mathcal{M})], \quad (9)$$

$$\text{BetP}_1(\theta_i) = \frac{1}{1 - m_1(\emptyset)} \sum_{\substack{A \subseteq \Theta \\ \theta_i \in A}} \frac{m_1(A)}{|A|}. \quad (10)$$

where $\mathcal{M} = \frac{1}{2}(\text{Bet}P_1 + \text{Bet}P_2)$ and $D_{KL}(\cdot\|\cdot)$ is the Kullback–Leibler divergence. The metric $d_{JS} \in [0, \ln 2]$ serves as a real-time measure of perceptual degradation.

The quantified uncertainty d_{JS} directly influences the pursuer's cost evaluation. An attenuated instantaneous cost during interference periods is introduced:

$$G^d(\mathbf{X}, U_P, U_E, \mu) = (1 - \zeta(t, \mu, d_{JS})) \cdot G(\mathbf{X}, U_P, U_E), \quad (11)$$

where $\mu \in [0, \mu_{\max}]$ is an interference intensity parameter. The attenuation function $\zeta(\cdot)$ is defined as:

$$\zeta(t, \mu, d_{JS}) = \frac{d_{JS} \cdot \mu}{\mathcal{K}_1 + d_{JS} \cdot \mu} \cdot [(t - t_w)(t_{w+1} - t)]^{\mathcal{K}_2} \cdot e^{-\mathcal{K}_3 \cdot (t - t_w)(t_{w+1} - t)}. \quad (12)$$

The parameters $\mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3 > 0$ shape the temporal profile of uncertainty, which satisfies $\zeta(t_w, \cdot) = \zeta(t_{w+1}, \cdot) = 0$ and is unimodal within the window $[t_w, t_{w+1}]$. The attenuation profile $\zeta(\cdot)$ is constructed to satisfy three practical requirements: (i) $\zeta(t_w, \cdot) = \zeta(t_{w+1}, \cdot) = 0$ so that boundary values are not distorted, (ii) $\zeta(\cdot)$ is unimodal within $[t_w, t_{w+1}]$ to avoid oscillatory re-weighting, and (iii) its amplitude scales with $d_{JS}(t_w)\mu(t_w)$ so that stronger degradation yields a more conservative cost shaping. A dedicated sensitivity study of $(\mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3)$ is reported in Section 6.7.

Evidence-theoretic modeling is adopted because the likelihood models required by fully Bayesian filters are often unavailable or unreliable under deliberate sensing denial, where observation statistics can be non-stationary and adversarially manipulated. Dempster–Shafer fusion supports multi-source evidence combination with explicit representation of ignorance and conflict, and the resulting pignistic probability offers an interpretable decision-ready belief over a finite maneuver-mode set. For the mode-cardinality considered here, the update and fusion steps remain lightweight and well suited to online window triggering.

3. Dynamic Time-Window Mechanism for Nash Equilibrium

The pursuit–evasion engagement is formulated in a dynamic time window framework, where the full horizon is partitioned into a sequence of subintervals. Window boundaries are generated online through event triggering rules. These rules are driven by changes in perceptual uncertainty, which is quantified using the Jensen Shannon divergence between the nominal belief and the interference corrupted belief. They are also driven by maneuver or strategy switches and by deviations in the state. With this partition, the differential game is expressed on a window by window basis, with a running cost that accounts for interference. Local Nash equilibrium strategies are then synthesized for each finite horizon subgame induced by a window. Accordingly, the time window update rule is introduced first. The triggering conditions are then stated, and an existence result is provided for a local Nash equilibrium in each dynamic window.

An overview of the proposed dynamic time-window mechanism is provided in Figure 2. The figure highlights the key coupling between perception quality and decision resolution: interference-corrupted beliefs increase the divergence indicator and thereby shorten the planning window, while the running cost is simultaneously adjusted to reflect uncertainty during high-degradation intervals. This design enables frequent policy refresh when sensing is unreliable and avoids unnecessary recomputation when perception is stable.




Figure 2. Dynamic time-window mechanism. Step (1) Acquire multi-source evidence and compute the belief $\text{BetP}(t)$ and degradation index $d_{JS}(t)$. Step (2) At the current boundary t_w , compute a candidate window length Δt_w via the smooth map $T(\cdot)$ in (17). Step (3) Within the candidate window, monitor event triggers: interference onset/termination (18) and (19), strategy-switch trigger (20), and state-deviation trigger (21). Step (4) Set the next boundary t_{w+1} by the composite rule (22) (with anti-chattering safeguards in Algorithm 1). Step (5) Solve the window-wise game and propagate the continuation value $V(X(t_{w+1}), t_{w+1})$ to ensure boundary consistency.

Algorithm 1 Unified dynamic time-window update with event triggers and anti-chattering

- 1: **Input:** current boundary t_w , state $X(t_w)$, belief $\text{BetP}(t_w)$, degradation $d_{JS}(t_w)$, intensity $\mu(t_w)$
- 2: Compute candidate length $\Delta t_w^{\text{cand}} \leftarrow T(d_{JS}(t_w), \mu(t_w), X(t_w))$ via (17)
- 3: Set $t_w^{\text{cand}} \leftarrow t_w + \Delta t_w^{\text{cand}}$
- 4: Initialize $t_{\text{next}} \leftarrow t_w^{\text{cand}}$
- 5: **Monitor triggers on** $(t_w, t_w^{\text{cand}}]$:
 Interference onset/termination using hysteresis ($d_{\text{on}}, d_{\text{off}}$)
 Strategy switch: $\|\text{BetP}(t) - \text{BetP}(t_w)\|_2 > \rho_{\text{switch}}$
 State deviation: $\|\bar{X}(t) - \bar{X}(t_w)\|_2 > \delta_{\text{state}}$
- 6: Let t_{trig} be the earliest time any trigger fires; if none, set $t_{\text{trig}} = +\infty$
- 7: **if** $t_{\text{trig}} < t_{\text{next}}$ **then**
- 8: $t_{\text{next}} \leftarrow t_{\text{trig}}$
- 9: **end if**
- 10: Enforce dwell time: $t_{w+1} \leftarrow \max\{t_{\text{next}}, t_w + \Delta t_{\text{dwell}}\}$
- 11: Enforce bounds: $t_{w+1} \leftarrow \min\{t_{w+1}, t_w + \Delta t_{\text{max}}\}$ and $\Delta t_w \leftarrow t_{w+1} - t_w \geq \Delta t_{\text{min}}$
- 12: **Output:** next boundary t_{w+1} and window length Δt_w

3.1. Time-Window Adaptation

The dynamic time-window mechanism divides the continuous pursuit–evasion game into sequential intervals denoted by $[t_w, t_{w+1}]$, where $w = 0, 1, \dots, W$. The length of each window $\Delta t_w = t_{w+1} - t_w$ is adaptively determined based on real-time perceptual conditions, particularly the level of uncertainty in the system’s state and the intensity of interference during the engagement.

The window adaptation function $\Delta t_w = T(d_{JS}(t_w), \mu(t_w), X(t_w))$ incorporates key factors that affect decision-making:

$$\Delta t_w = T(d_{JS}(t_w), \mu(t_w), X(t_w)), \quad (13)$$

where $d_{JS}(t_w)$ is the Jensen–Shannon divergence at time t_w , representing the perceptual uncertainty due to incomplete information, $\mu(t_w)$ is the interference intensity at time

t_w , reflecting the level of active countermeasures employed by the evader, $X(t_w)$ is the state of the system at time t_w , including the relative position and velocity of the pursuer and evader.

We first specify the qualitative requirements that the window-length map $T(\cdot)$ must satisfy (monotonicity, boundedness, and smoothness), which are independent of any particular functional choice. A concrete realization is then selected to meet these requirements while avoiding discontinuous window-length jumps; this leads to the sigmoidal law in (17), which provides a smooth interpolation between $\Delta t_{\min}$ and $\Delta t_{\max}$.

The adaptation function is designed to satisfy the following properties:

$$\frac{\partial T}{\partial d_{JS}} < 0, \quad (14)$$

indicating that the window length decreases as uncertainty increases, and

$$\frac{\partial T}{\partial \mu} < 0, \quad (15)$$

implying that higher interference leads to shorter decision windows. Additionally, the boundaries of the window length is set as:

$$T(0, 0, X) = \Delta t_{\max} \quad \text{and} \quad T(d_{\max}^{JS}, \mu_{\max}, X) = \Delta t_{\min}, \quad (16)$$

ensuring that the window length is bounded within the limits defined by the maximum and minimum tolerances for decision-making.

A sigmoidal adaptation law is proposed for smooth transitions between window lengths:

$$\Delta t_w = \Delta t_{\min} + \frac{\Delta t_{\max} - \Delta t_{\min}}{1 + \exp(\beta_1 \cdot d_{JS}(t_w) + \beta_2 \cdot \mu(t_w) - \beta_0)}, \quad (17)$$

where $\beta_0, \beta_1, \beta_2 > 0$ are sensitivity parameters that determine how the window length responds to uncertainty and interference.

3.2. Window Triggering Conditions

The start time of the w -th time window is denoted by t_w . The next boundary t_{w+1} is determined in an event-triggered manner according to changes in the perceptual environment. Specifically, four types of triggering conditions are considered: (i) interference onset, (ii) interference termination, (iii) strategy switch, and (iiii) state deviation. To avoid numerical sensitivity (e.g., $d_{JS}(t)$ rarely equals exactly zero in practice), we introduce a pair of interference thresholds d_{on} and d_{off} with $d_{\text{on}} > d_{\text{off}} \geq 0$ to realize a hysteresis switching mechanism. Throughout, we adopt the convention $\inf \emptyset = +\infty$.

- **Interference Onset:** If the system is currently regarded as non-interfered, i.e., $d_{JS}(t_w) \geq d_{\text{on}}$, then the candidate triggering time for interference onset is

$$t_{\text{on}} \triangleq \inf \left\{ t \in (t_w, t_w + \Delta t_{\max}] \mid d_{JS}(t) \geq d_{\text{on}} \right\}, \quad (18)$$

otherwise we set $t_{\text{on}} = +\infty$.

- **Interference Termination:** If the system is currently regarded as interfered, i.e., $d_{JS}(t_w) > d_{\text{off}}$, then the candidate triggering time for interference termination is

$$t_{\text{off}} \triangleq \inf \left\{ t \in (t_w, t_w + \Delta t_{\max}] \mid d_{JS}(t) \leq d_{\text{off}} \right\}, \quad (19)$$

otherwise we set $t_{\text{off}} = +\infty$.

- **Strategy Switch Detection:** A new window is triggered when the evader's strategy changes significantly, measured by the change in the belief distribution:

$$t_{\text{strategy}} \triangleq \inf \left\{ t \in (t_w, t_w + \Delta t_{\text{max}}] \mid \|\text{BetP}(t) - \text{BetP}(t_w)\|_2 \geq \rho_{\text{switch}} \right\}. \quad (20)$$

where the scalar $\rho_{\text{switch}} > 0$ is a user-specified threshold that determines whether the evader's strategy has changed significantly.

- **State-Based Trigger:** A new window is triggered when the system state deviates sufficiently from the state at t_w . Since $X = [r^\top, v^\top]^\top$ mixes different physical units (m and m/s), we use a normalized weighted norm:

$$t_{\text{state}} \triangleq \inf \left\{ t \in (t_w, t_w + \Delta t_{\text{max}}] \mid \|S(X(t) - X(t_w))\|_2 \geq \delta_{\text{state}} \right\}, \quad (21)$$

where S is a diagonal scaling matrix, e.g., $S = \text{diag}(1/r_0, 1/r_0, 1/r_0, 1/v_0, 1/v_0, 1/v_0)$ with characteristic scales r_0, v_0 and the scalar $\delta_{\text{state}} > 0$ is a user-specified threshold that declares a significant state change.

Finally, the composite triggering rule selects the earliest event time, while ensuring that the window length never exceeds Δt_{max} even if no triggering event occurs:

$$t_{w+1} = \min \{ t_{\text{on}}, t_{\text{off}}, t_{\text{strategy}}, t_{\text{state}}, t_w + \Delta t_{\text{max}} \}. \quad (22)$$

To avoid chattering under noisy belief updates, we implement three safeguards. First, a hysteresis pair $(d_{\text{on}}, d_{\text{off}})$ with $d_{\text{on}} > d_{\text{off}}$ is used for interference onset/termination, so that small oscillations of $d_{\text{JS}}(t)$ around a single threshold do not trigger rapid switching. Second, a minimum dwell time Δt_{dwell} is enforced between consecutive boundaries. Third, the trigger signals $\text{BetP}(t)$ and $X(t)$ are evaluated using a short moving average (or a first-order low-pass filter) before computing $\|\text{BetP}(t) - \text{BetP}(t_w)\|_2$ and $\|X(t) - X(t_w)\|_2$. These safeguards are summarized in Algorithm 1.

Figure 3 summarizes the four event-triggering conditions used to generate window boundaries, including interference onset/termination (implemented with hysteresis thresholds to avoid chattering), strategy-switch detection based on belief variation, and a state-deviation trigger based on a normalized weighted norm. The next boundary is selected as the earliest among these candidate event times, which guarantees $\Delta t_w \leq \Delta t_{\text{max}}$ and makes the partition responsive to both information changes and dynamical deviations.




Figure 3. Four Event-Triggered Dynamic Time-Window Switching Mechanisms.

3.3. Existence of a Local Nash Equilibrium in Each Dynamic Time Window

Section 3.2 specifies an event-triggered mechanism that generates a time window $\{[t_w, t_{w+1}]\}_{w=0}^{W-1}$ with $t_{w+1} \in (t_w, t_w + \Delta t_{\max}]$. In order for the window-wise HJBI synthesis in Section 3.2 to be well-posed, it remains to justify that each induced finite-horizon subgame admits a Nash equilibrium (saddle-point) solution. This subsection provides a rigorous existence result.

Fix a window index w and consider the pursuit–evasion game restricted to the time interval $[t_w, t_{w+1}]$, with state $X(t) \in \mathbb{R}^n$ and admissible control processes $U_P(\cdot) \in \mathcal{U}_P([t_w, t_{w+1}])$ and $U_E(\cdot) \in \mathcal{U}_E([t_w, t_{w+1}])$. Following the game tuple in Equation (5), the dynamics are written as

$$\dot{X}(t) = f(X(t), U_P(t), U_E(t)), \quad t \in [t_w, t_{w+1}], \quad (23)$$

with the initial condition $X(t_w) = x$. In our spacecraft relative-motion setting, (23) is induced by the CW/LVLH state-space model (2) by taking the relative control as $U(t) = U_P(t) - U_E(t)$, i.e., $\dot{X} = AX + B(U_P - U_E)$.

The cost functional within each window is defined as:

$$J_w(X, U_P, U_E) = \int_{t_w}^{t_{w+1}} G_w(X(\tau), U_P(\tau), U_E(\tau), \mu(\tau)) d\tau + V(X(t_{w+1}), t_{w+1}), \quad (24)$$

where $G_w(X, U_P, U_E, \mu)$ is the window-specific instantaneous cost, modified by the perceptual uncertainty and interference levels. $V(X(t_{w+1}), t_{w+1})$ is the value function at the window boundary, which encodes the terminal cost at the end of the window. In implementation, the continuation value $V(X(t_{w+1}), t_{w+1})$ is approximated numerically by warm-starting the next-window solver from the boundary state and evaluating the window value function on a short look-ahead horizon; details of this approximation and its interaction with the VI/ADMM solver are provided in Section 5 and validated in Section 6.

During periods of high interference, the instantaneous cost is attenuated according to the uncertainty function $\zeta(t, \mu, d_{JS})$, which reduces the influence of control efforts on the overall cost. Thus, the cost function within each window becomes:

$$G_w(X, U_P, U_E) = [1 - \zeta(t, \mu, d_{JS})] \cdot G(X, U_P, U_E) + \zeta(t, \mu, d_{JS}) \cdot R(X, U_P), \quad (25)$$

where $R(X, U_P)$ is a robustness regularization term specified as a conservative constraint to penalize overly aggressive closing actions under degraded perception:

$$R(X, U_P) = \alpha_r r^2 + \alpha_v \|v\|_2^2 + \beta_r \|U_P\|_2^2, \quad (26)$$

where $X = [r^\top \ v^\top]^\top$ stacks the relative position and velocity, and $\alpha_v > 0$ discourages large relative speed when sensing is unreliable. The weights $(\alpha_r, \alpha_v, \beta_r)$ are chosen no smaller than the nominal weights to enforce conservativeness during high-degradation intervals.

The continuation value $V(\cdot, t_{w+1})$ plays the role of a terminal cost at the window boundary and is provided by the backward recursion in Section 4 (or by a terminal condition at T).

Let $U_P \subset \mathbb{R}^{m_P}$ and $U_E \subset \mathbb{R}^{m_E}$ be the compact control constraint sets specified in Section 2.2. Define $\mathcal{U}_P([t_w, t_{w+1}])$ and $\mathcal{U}_E([t_w, t_{w+1}])$ as the sets of measurable, essentially bounded functions on $[t_w, t_{w+1}]$ taking values in U_P and U_E , respectively. We impose the following standard conditions.

Assumption 1. f is continuous in (X, U_P, U_E) and locally Lipschitz in X uniformly over $(U_P, U_E) \in U_P \times U_E$.

Assumption 2. $\mu(t)$, $d_{JS}(t)$ are measurable and bounded on $[t_w, t_{w+1}]$, and $\xi(t, \mu, d_{JS})$ is measurable with $0 \leq \xi \leq 1$.

Assumption 3. $G(\cdot)$ and $R(\cdot)$ are continuous in (X, U_P, U_E) and (X, U_P) , respectively, and $V(\cdot, t_{w+1})$ is continuous and bounded from below.

Assumption 4. The window Hamiltonian satisfies the Isaacs minimax equality (defined below). This condition holds, for instance, under standard convexification via relaxed/mixed controls, or under additional structural properties that render the static min–max and max–min equivalent almost everywhere.

Remark 2 (Practical limitations of Isaacs condition). Assumption 4 is a sufficient condition that may fail in practice when control constraints induce nonconvex Hamiltonians, when the optimal feedback is nonsmooth, or when actuation/estimation delays break the instantaneous min–max interchange. In spacecraft proximity operations, such violations can occur under hard-thrust saturation, safety-driven nonconvex state constraints, or highly anisotropic uncertainty. In this work, the Isaacs condition is adopted to support the HJBI characterization; numerically, we additionally monitor the saddle residual $|\min_{U_P} \max_{U_E} \mathcal{H} - \max_{U_E} \min_{U_P} \mathcal{H}|$ within each window to ensure consistency of the computed equilibria.

For $p \in \mathbb{R}^n$, define the window Hamiltonian

$$\mathcal{H}_w(X, p, t) \triangleq \min_{u_P \in U_P} \max_{u_E \in U_E} \left\{ p^\top f(X, u_P, u_E) + G_w(X, u_P, u_E, \mu(t)) \right\}, \quad (27)$$

and, by Assumption 4, equivalently

$$\mathcal{H}_w(X, p, t) = \max_{u_E \in U_E} \min_{u_P \in U_P} \left\{ p^\top f(X, u_P, u_E) + G_w(X, u_P, u_E, \mu(t)) \right\}. \quad (28)$$

Theorem 1. Consider any window $[t_w, t_{w+1}]$ generated by the triggering rule in Section 3.2, so that $t_{w+1} \in (t_w, t_w + \Delta t_{\max}]$. Under Assumptions 1–4, the finite-horizon zero-sum differential game (5) admits a (window-wise) value function $V_w : \mathbb{R}^n \times [t_w, t_{w+1}] \rightarrow \mathbb{R}$ such that, for all $(x, t) \in \mathbb{R}^n \times [t_w, t_{w+1}]$,

$$\begin{aligned} V_w(x, t) &= \inf_{U_P \in \mathcal{U}_P([t, t_{w+1}])} \sup_{U_E \in \mathcal{U}_E([t, t_{w+1}])} J_w(x; U_P, U_E) \\ &= \sup_{U_E \in \mathcal{U}_E([t, t_{w+1}])} \inf_{U_P \in \mathcal{U}_P([t, t_{w+1}])} J_w(x; U_P, U_E). \end{aligned} \quad (29)$$

Moreover, V_w is the (unique) viscosity solution of the window-wise HJBI equation

$$-\partial_t V_w(x, t) = \mathcal{H}_w(x, \nabla_x V_w(x, t), t), \quad t \in [t_w, t_{w+1}], \quad (30)$$

with terminal condition $V_w(x, t_{w+1}) = V(x, t_{w+1})$. Finally, there exist admissible saddle controls $(U_{P,w}^*, U_{E,w}^*)$ on $[t_w, t_{w+1}]$ (the local Nash equilibrium on the window) such that

$$J_w(x; U_{P,w}^*, U_E) \leq J_w(x; U_{P,w}^*, U_{E,w}^*) \leq J_w(x; U_P, U_{E,w}^*), \quad \forall U_P \in \mathcal{U}_P, \forall U_E \in \mathcal{U}_E. \quad (31)$$

Proof. Step 1 (Well-posedness on a bounded window). By Section 3.3, $t_{w+1} - t_w \leq \Delta t_{\max} < \infty$. Fix any admissible (U_P, U_E) . Assumption 1 implies existence and uniqueness of an absolutely continuous trajectory $X(\cdot)$ solving (23) on $[t_w, t_{w+1}]$ for each initial state $X(t_w) = x$.

Step 2 (Well-defined payoff). Assumptions 2–3 guarantee that $\mu(\cdot)$, $d_{JS}(\cdot)$ and $\xi(\cdot)$ are bounded and measurable on the window and that $G_w(\cdot)$ is measurable in time and continuous in (X, U_P, U_E) . Therefore the integral term in (24) is well-defined, and adding the continuous terminal term $V(X(t_{w+1}), t_{w+1})$ yields a finite payoff $J_w(x; U_P, U_E)$.

Step 3 (Lower/upper values and Isaacs equality). Define the upper and lower values on $[t, t_{w+1}]$:

$$\begin{aligned}\bar{V}_w(x, t) &\triangleq \inf_{U_P \in \mathcal{U}_P([t, t_{w+1}])} \sup_{U_E \in \mathcal{U}_E([t, t_{w+1}])} J_w(x; U_P, U_E), \\ \underline{V}_w(x, t) &\triangleq \sup_{U_E \in \mathcal{U}_E([t, t_{w+1}])} \inf_{U_P \in \mathcal{U}_P([t, t_{w+1}])} J_w(x; U_P, U_E).\end{aligned}\quad (32)$$

Standard dynamic programming arguments for finite-horizon differential games on compact time intervals imply that $\bar{V}_w$ and $\underline{V}_w$ satisfy the upper and lower Isaacs equations in the viscosity sense, sharing the same terminal condition $V(\cdot, t_{w+1})$ at t_{w+1} . Assumption 4 ensures that the upper and lower Hamiltonians coincide; hence, $\bar{V}_w \equiv \underline{V}_w$ and (29) holds with the common value denoted V_w .

Step 4 (Existence of a saddle-point on the window). Because U_P and U_E are compact and the integrand in (27) is continuous in (u_P, u_E) for fixed (X, p, t) , the extrema in the Hamiltonian are attained. Using the measurable dependence on t (Assumption 2) and standard measurable selection arguments, one can select admissible controls achieving the saddle structure induced by the Isaacs equality, which yields the saddle inequalities (31).

Step 5 (Effect of the dynamic triggering mechanism). The boundary t_{w+1} is generated by the event conditions in Section 3.3 from measurable signals (state, belief, and interference indicators), and is always bounded by $t_w + \Delta t_{\max}$. Conditioning on the information available at t_w , the ensuing subgame is a standard finite-horizon differential game posed on the deterministic interval $[t_w, t_{w+1}]$, so Steps 1–4 apply to every realized window. Therefore, a local Nash equilibrium exists in each time window produced by the dynamic time-window mechanism. $\square$

Remark 3. Theorem 1 validates the key assumption used in Section 4, namely that every window-wise subproblem admits a Nash equilibrium (saddle-point) solution. Together with the consistency conditions in Section 4.2, this enables the concatenation argument establishing global game-theoretic consistency over $[0, T]$.

4. Global Nash Equilibrium Consistency and Robust Optimization

Section 3 provides a window-wise equilibrium construction under dynamic time windows. This section establishes (i) when the concatenation of window equilibria yields a full-horizon equilibrium (global consistency) and (ii) how to embed each window subproblem into a variational distributionally robust optimization (VDRO) layer to hedge against ambiguity in interference intensity.

4.1. Induced Global Game Under Dynamic Time Windows

Let the horizon $[0, T]$ be partitioned online into consecutive windows $\{[t_w, t_{w+1}]\}_{w=0}^W$, where $t_0 = 0$ and $t_{W+1} = T$. Denote, using $U_{P,w}$ and $U_{E,w}$, the restrictions of the pursuer/evader controls onto $[t_w, t_{w+1}]$.

Following the window-wise synthesis in Section 3, let $(U_{P,w}^*, U_{E,w}^*)$ be a local Nash equilibrium (saddle) on window w computed given the boundary information at t_w . The concatenated strategy pair is defined as

$$(U_P^*, U_E^*) \triangleq \{(U_{P,0}^*, U_{E,0}^*), \dots, (U_{P,W}^*, U_{E,W}^*)\}. \quad (33)$$

For notational clarity, let J_g denote the full-horizon payoff, which coincides with the cost function J in Equation (6) when evaluated over $[0, T]$. Within each window, the corresponding objective follows the window payoff definition in Equation (24), with the terminal term taken as a continuation value at the window boundary (defined below). This design enforces dynamic-programming compatibility across window interfaces.

4.2. Global Continuation Value Recursion and Consistency Theorem

Define the global continuation value $V_g(\cdot, \cdot)$ via backward recursion over the realized dynamic windows. Specifically,

$$V_g(X, t) \triangleq \begin{cases} H(X(T)), & t = T, \\ \min_{U_{P,w} \in \mathcal{U}_P} \max_{U_{E,w} \in \mathcal{U}_E} \left[\int_t^{t_{w+1}} G_w(X(\tau), U_{P,w}(\tau), U_{E,w}(\tau), \mu(\tau)) d\tau + V_g(X(t_{w+1}), t_{w+1}) \right], & t \in [t_w, t_{w+1}), \end{cases} \quad (34)$$

where w is the unique index satisfying $t \in [t_w, t_{w+1})$, and $\mathcal{U}_P, \mathcal{U}_E$ denote the admissible strategy classes used throughout (e.g., measurable policies valued in compact convex sets).

Theorem 2 (Global Nash equilibrium consistency). *Suppose that, for each window $[t_w, t_{w+1}]$, there exists a local Nash equilibrium $(U_{P,w}^*, U_{E,w}^*)$ with respect to the window objective using the continuation value $V_g(X(t_{w+1}), t_{w+1})$. Assume further that the following structural properties hold:*

(C1) *State continuity at window boundaries. The physical state trajectory is continuous at each boundary:*

$$\lim_{\varepsilon \rightarrow 0^+} X(t_w - \varepsilon) = \lim_{\varepsilon \rightarrow 0^+} X(t_w + \varepsilon) = X(t_w), \quad \forall w \in \{1, \dots, W\}. \quad (35)$$

(C2) *Additive decomposition of the global cost under the partition. For any admissible (U_P, U_E) , the global objective admits the window-additive form*

$$J_g(X_0, U_P, U_E) = \sum_{w=0}^W \int_{t_w}^{t_{w+1}} G_w(X(t), U_{P,w}(t), U_{E,w}(t), \mu(t)) dt + H(X(T)), \quad (36)$$

where $U_{P,w}, U_{E,w}$ are the restrictions of U_P, U_E onto $[t_w, t_{w+1}]$.

(C3) *Markovian (memoryless) structure of local equilibria. Each local equilibrium strategy depends on the current state (and time) only:*

$$U_{P,w}^*(t) = \pi_{P,w}^*(X(t), t), \quad U_{E,w}^*(t) = \pi_{E,w}^*(X(t), t), \quad t \in [t_w, t_{w+1}]. \quad (37)$$

Then the concatenated pair (U_P^*, U_E^*) in (33) constitutes a Nash equilibrium for the full-horizon game on $[0, T]$, i.e.,

$$J_g(X_0, U_P^*, U_E) \leq J_g(X_0, U_P^*, U_E^*) \leq J_g(X_0, U_P, U_E^*), \quad \forall U_P \in \mathcal{U}_P, \forall U_E \in \mathcal{U}_E. \quad (38)$$

Proof. We prove (38) by a one-step (window-local) deviation argument and dynamic programming compatibility induced by the continuation value.

Step 1: window-wise saddle property with continuation value. By the definition of a local Nash equilibrium on window $[t_w, t_{w+1}]$ for the payoff function, we have, for each fixed boundary state $X(t_w)$,

$$J_w(X(t_w), U_{P,w}^*, U_{E,w}) \leq J_w(X(t_w), U_{P,w}^*, U_{E,w}^*) \leq J_w(X(t_w), U_{P,w}, U_{E,w}^*), \quad \forall U_{P,w}, U_{E,w}. \quad (39)$$

Step 2: unilateral deviation of the pursuer. Fix any pursuer strategy $U_P \in \mathcal{U}_P$ and consider the pair (U_P, U_E^*) . If $U_P \equiv U_P^*$, the right inequality in (38) is trivial. Otherwise, let $w^\dagger$ be the *first* window index where U_P differs from U_P^* on a set of nonzero measure. By construction, the strategies coincide on all earlier windows $[t_0, t_1], \dots, [t_{w^\dagger-1}, t_{w^\dagger}]$. Under (C1), the boundary state $X(t_{w^\dagger})$ is well-defined and identical for the two pairs (U_P, U_E^*) and (U_P^*, U_E^*) up to time $t_{w^\dagger}$. Moreover, by (C3), the continuation policy $(U_{P,k}^*, U_{E,k}^*)$ on later windows $k > w^\dagger$ depends on the past only through the boundary state; hence, once the boundary state is fixed, the comparison between the deviating and equilibrium controls on window $w^\dagger$ can be carried out using the same continuation value $V_g(X(t_{w^\dagger+1}), t_{w^\dagger+1})$.

Applying the right inequality of (39) with $w = w^\dagger$ and $U_{E,w^\dagger} = U_{E,w^\dagger}^*$, we obtain

$$\begin{aligned} & \int_{t_{w^\dagger}}^{t_{w^\dagger+1}} G_w(X(t), U_{P,w^\dagger}(t), U_{E,w^\dagger}^*(t), \mu(t)) dt + V_g(X(t_{w^\dagger+1}), t_{w^\dagger+1}) \\ & \geq \int_{t_{w^\dagger}}^{t_{w^\dagger+1}} G_w(X(t), U_{P,w^\dagger}^*(t), U_{E,w^\dagger}^*(t), \mu(t)) dt + V_g(X(t_{w^\dagger+1}), t_{w^\dagger+1}). \end{aligned} \quad (40)$$

Canceling the common continuation term yields a non-improvement inequality for the deviating window.

Now invoke (C2) to decompose the full-horizon cost into the sum of window contributions. All windows $k < w^\dagger$ are unchanged by the choice of $w^\dagger$, and windows $k > w^\dagger$ are evaluated under the equilibrium continuation. Summing (40) with the unchanged window costs implies

$$J_g(X_0, U_P, U_E^*) \geq J_g(X_0, U_P^*, U_E^*), \quad \forall U_P \in \mathcal{U}_P, \quad (41)$$

which is exactly the right inequality in (38).

Step 3: unilateral deviation of the evader. The proof of the left inequality is symmetric. Fix any $U_E \in \mathcal{U}_E$ and define the first window where U_E deviates from U_E^* . Using the left inequality of (39) on that window, together with (C1)–(C3), and summing via (C2), yields

$$J_g(X_0, U_P^*, U_E) \leq J_g(X_0, U_P^*, U_E^*), \quad \forall U_E \in \mathcal{U}_E. \quad (42)$$

Combining (41) and (42) establishes (38). $\square$

The assumptions (C1)–(C3) are introduced to guarantee dynamic-programming compatibility under window partitioning. In particular, restricting to local Markov feedback policies ensures that the boundary continuation value depends on $(X(t_{w+1}), t_{w+1})$ rather than the full history. When history-dependent behaviors are essential, the same analysis can be extended by augmenting the state with a finite-memory summary (e.g., filtered belief statistics), at the expense of increased computational load. Thus, Theorem 2 should be interpreted as a rigorous consistency guarantee within the practically relevant class of online-computable Markovian strategies.

In this work, the window-partition rule is treated as an exogenous system mechanism: the triggering logic and thresholds are fixed, common knowledge to both players and are not strategic decision variables. Although the trigger signals depend on the realized state/belief trajectories (and thus indirectly on controls), the equilibrium concept is defined for the resulting hybrid time-partitioned game under this prescribed partition mechanism. Endogenizing the switching times as additional strategic variables leads to a different class of hybrid differential games, which is beyond the present scope and is left for future investigation.

Remark 4 (Interpretation of the assumptions). *Condition (C1) is physical for continuous-time spacecraft dynamics; (C2) holds when the window running cost G_w is defined consistently with*

the global integrand and the terminal cost appears only once at T ; and (C3) corresponds to the Markovian equilibrium class induced by HJBI-type characterizations in Section 3. Together with the continuation value V_g , these properties ensure that the window-wise equilibrium is subgame consistent, and thus, its concatenation forms a full-horizon Nash equilibrium.

4.3. Window-Wise VDRO Embedding Under Interference Ambiguity

Within each window, the interference intensity μ is uncertain and potentially adversarial. Let $\mathbb{P}$ denote the (unknown) distribution of μ supported on $\Xi = [0, \mu_{\max}]$. We adopt a moment-based ambiguity set

$$\mathcal{P}_w \triangleq \left\{ \mathbb{P} \in \mathcal{M}_+(\Xi) \mid \mathbb{E}_{\mathbb{P}}[\mu] = \mu_0(t_w), \mathbb{E}_{\mathbb{P}}[\mu^2] \leq \mu_0^2(t_w) + \sigma^2(t_w), \text{supp}(\mathbb{P}) \subseteq [0, \mu_{\max}] \right\}. \quad (43)$$

The window-level moments $\mu_0(t_w)$ and $\sigma^2(t_w)$ are estimated online from multi-modal perception/evidence fusion over a short history horizon and can be made adaptive to perceptual degradation (e.g., larger uncertainty $\Rightarrow$ larger ambiguity radius).

We adopt a moment-based ambiguity set because only low-order statistics of the interference intensity can be reliably estimated online over short windows, especially under sensing denial. This choice yields a tractable dual representation with a closed-form worst-case density structure, enabling efficient inner updates. While Wasserstein-based sets can provide finer geometry of distributional deviation, they typically require either larger sample sizes or more expensive inner maximization, which is less compatible with strict per-window real-time budgets considered here.

Given a fixed evader strategy within an inner iteration, the pursuer solves a VDRO best-response:

$$\min_{U_{P,w} \in \mathcal{U}_P} \max_{\mathbb{P} \in \mathcal{P}_w} \mathbb{E}_{\mathbb{P}} \left[\int_{t_w}^{t_{w+1}} G_w(X(t), U_{P,w}(t), U_{E,w}(t), \mu(t)) dt + H(X(t_{w+1})) \right]. \quad (44)$$

A dual representation with multipliers associated with the moment constraints yields a worst-case density of (truncated) exponential-family form on $[0, \mu_{\max}]$, which provides a computationally convenient update for $\mathbb{P}$. In practice, we alternate between (i) control optimization under the current worst-case distribution and (ii) distribution update within $\mathcal{P}_w$:

$$U_{P,w}^{(k+1)} = \arg \min_{U_{P,w} \in \mathcal{U}_P} \mathbb{E}_{\mathbb{P}^{(k)}} [J_w(U_{P,w}, \mu)] \quad (45)$$

$$\mathbb{P}^{(k+1)} = \arg \max_{\mathbb{P} \in \mathcal{P}_w} \mathbb{E}_{\mathbb{P}} [J_w(U_{P,w}^{(k+1)}, \mu)] \quad (46)$$

Under standard convexity/regularity conditions induced by the quadratic running costs and the regularization parameter η , the alternating updates in (45) and (46) are expected to contract toward a stationary saddle of the window subproblem. In real-time use, we employ warm starts across consecutive windows and terminate each inner loop after a fixed iteration budget (or when both the primal residual and dual residual fall below prescribed tolerances), which provides predictable latency per window. Empirically, this early-stopping strategy preserves solution quality while ensuring online feasibility.

In summary, this section bridges the dynamic time-window equilibrium construction with a VDRO layer that accounts for distributional ambiguity in the interference intensity. The global consistency result ensures that window-wise saddle solutions can be concatenated without violating the full-horizon equilibrium semantics, while the moment-based ambiguity set provides a principled robustness mechanism within each window.

Computationally, the resulting robust window subproblem is a coupled min–max program. To enable efficient real-time solution under limited window-level iterations,

the next section derives its KKT conditions and reformulates them as a finite-dimensional variational inequality, which naturally motivates the accelerated operator-splitting (ADMM-type) algorithm.

5. Variational Inequality Reformulation of KKT Conditions and Accelerated ADMM Solver

This section develops an efficient solver for the window-wise distributionally robust subproblems introduced in Section 4.3. We first derive a well-posed KKT system by employing an entropic smoothing of the inner worst-case distribution update, which yields a unique Gibbs-form worst-case density and a finite-dimensional parametrization. We then reformulate the resulting optimality system as a variational inequality (VI) and design an ADMM-type operator-splitting scheme with optional acceleration for real-time, window-level computation.

5.1. KKT Conditions for Window-Wise VDRO Subproblem

Before presenting the detailed derivations, we summarize the solver pipeline. For each window, the interference intensity is modeled by a moment-constrained ambiguity set, yielding a VDRO min–max problem. We then derive the associated KKT system of the inner distribution update and reformulate it as a finite-dimensional VI, whose structured monotone-inclusion form admits an efficient operator-splitting implementation. Finally, an accelerated ADMM-type iteration is applied per window with warm starts to meet online computational budgets.

Consider a fixed time window $[t_w, t_{w+1}]$ and abbreviate the corresponding window objective by $J_w(U_P, \mu)$, where $U_P \in \mathcal{U}_P$ denotes the pursuer control strategy restricted to the window and $\mu \in [0, \mu_{\max}]$ is the (uncertain) interference intensity. Following (44), the pursuer’s VDRO best-response (with the evader treated as fixed within one inner-loop step) is

$$\min_{U_P \in \mathcal{U}_P} \max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[J_w(U_P, \mu)], \quad (47)$$

where $\mathcal{P}$ is the moment-based ambiguity set in (43), i.e., $\mathbb{E}_{\mathbb{P}}[\mu] = \mu_0$, $\mathbb{E}_{\mathbb{P}}[\mu^2] \leq \mu_0^2 + \sigma^2$, and $\text{supp}(\mathbb{P}) \subseteq [0, \mu_{\max}]$.

To obtain a single-valued mapping from finite-dimensional dual variables to the worst-case distribution (and hence a stable VI operator), we adopt an entropy-regularized variant of the inner maximization:

$$\max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[J_w(U_P, \mu)] + \eta_{\text{reg}} \mathcal{H}(\mathbb{P}), \quad \eta_{\text{reg}} > 0, \quad (48)$$

where $\mathcal{H}(\mathbb{P})$ denotes the (differential) entropy when $\mathbb{P}$ admits a density on $[0, \mu_{\max}]$. This smoothing is used as an algorithmic regularizer; the baseline (unregularized) worst-case update is recovered in the limit of vanishing smoothing, i.e., as the smoothing parameter becomes arbitrarily small. The parameter $\eta_{\text{reg}} > 0$ plays the role of a temperature (entropy-regularization strength) in the truncated exponential-family form (48). Smaller η_{reg} yields a more concentrated density $p^*(\mu)$ that places most mass near the most adverse interference intensities (thus approaching an extreme worst-case distribution), while larger η_{reg} produces a smoother density that improves numerical conditioning and stabilizes the inner update. In our implementation, η_{reg} is selected to balance worst-case sharpness and solver robustness.

Proposition 1 (Limit as $\eta_{\text{reg}} \downarrow 0$). Let $\Phi(\mu) \triangleq J_w(\mu) + \lambda_1 \mu + \lambda_2 \mu^2$ be continuous on $[0, \mu_{\max}]$, and define $p_{\eta_{\text{reg}}}(\mu) \propto \exp(-\Phi(\mu)/\eta_{\text{reg}}) \mathbf{1}_{[0, \mu_{\max}]}(\mu)$. Then, as $\eta_{\text{reg}} \downarrow 0$, (i) $p_{\eta_{\text{reg}}}$ converges weakly

to a probability measure supported on $\arg \min_{\mu \in [0, \mu_{\max}]} \Phi(\mu)$, and (ii) the log-partition value satisfies the Laplace principle

$$-\eta_{\text{reg}} \log \int_0^{\mu_{\max}} \exp\left(-\frac{\Phi(\nu)}{\eta_{\text{reg}}}\right) d\nu \longrightarrow \min_{\mu \in [0, \mu_{\max}]} \Phi(\mu). \quad (49)$$

Consequently, the regularized inner update approaches an extreme (unregularized) worst-case behavior in the small- η_{reg} limit.

Proof. The result follows from the Laplace principle for exponential integrals on compact domains. $\square$

Introduce Lagrange multipliers $\lambda_1 \in \mathbb{R}$ (mean equality) and $\lambda_2 \geq 0$ (second-moment upper bound). For clarity, we write the second-moment constraint in the equivalent form

$$\mu_0^2 + \sigma^2 - \mathbb{E}_{\mathbb{P}}[\mu^2] \geq 0, \quad (50)$$

so that its multiplier naturally satisfies $\lambda_2 \geq 0$.

The (distributional) Lagrangian of the entropy-regularized inner problem (48) can be expressed (up to the normalization constraint) as

$$\mathcal{L}(\mathbb{P}, U_P, \lambda_1, \lambda_2) = \mathbb{E}_{\mathbb{P}}[J_w(U_P, \mu)] + \lambda_1(\mathbb{E}_{\mathbb{P}}[\mu] - \mu_0) + \lambda_2(\mu_0^2 + \sigma^2 - \mathbb{E}_{\mathbb{P}}[\mu^2]) + \eta_{\text{reg}} \mathcal{H}(\mathbb{P}). \quad (51)$$

Let $\mathbb{P}^*$ denote a worst-case distribution and $(U_P^*, \lambda_1^*, \lambda_2^*)$ denote corresponding saddle variables. The KKT conditions consist of:

Primal feasibility.

$$\mathbb{E}_{\mathbb{P}^*}[\mu] = \mu_0, \quad \mathbb{E}_{\mathbb{P}^*}[\mu^2] \leq \mu_0^2 + \sigma^2, \quad \text{supp}(\mathbb{P}^*) \subseteq [0, \mu_{\max}]. \quad (52)$$

Dual feasibility.

$$\lambda_2^* \geq 0. \quad (53)$$

Complementary slackness.

$$\lambda_2^*(\mu_0^2 + \sigma^2 - \mathbb{E}_{\mathbb{P}^*}[\mu^2]) = 0. \quad (54)$$

Define the regularized worst-case value

$$\Phi_w(U_P) \triangleq \max_{\mathbb{P} \in \mathcal{P}} (\mathbb{E}_{\mathbb{P}}[J_w(U_P, \mu)] + \eta_{\text{reg}} \mathcal{H}(\mathbb{P})). \quad (55)$$

Under standard regularity (so that Danskin's theorem applies), stationarity of the outer minimization can be written as

$$\nabla_{U_P} \Phi_w(U_P) \Big|_{U_P=U_P^*} = 0, \quad \text{equivalently} \quad \nabla_{U_P} \mathbb{E}_{\mathbb{P}^*}[J_w(U_P, \mu)] \Big|_{U_P=U_P^*} = 0. \quad (56)$$

The entropy regularization in (48) yields a truncated exponential-family density on $[0, \mu_{\max}]$:

$$p^*(\mu) = \frac{\exp\left(\frac{1}{\eta} [J_w(U_P^*, \mu) + \lambda_1^* \mu - \lambda_2^* \mu^2]\right)}{\int_0^{\mu_{\max}} \exp\left(\frac{1}{\eta} [J_w(U_P^*, \nu) + \lambda_1^* \nu - \lambda_2^* \nu^2]\right) d\nu} \mathbf{1}_{[0, \mu_{\max}]}(\mu), \quad (57)$$

where $\eta > 0$ controls regularity and numerical stability.

5.2. Finite-Dimensional Variational Inequality Reformulation

To avoid direct optimization over the infinite-dimensional distribution space, we embed $\mathbb{P}^*$ into a finite-dimensional mapping via (57). Define the stacked primal–dual variable

$$z \triangleq (U_P, \lambda_1, \lambda_2), \quad (58)$$

and the feasible set

$$\mathcal{Z} \triangleq \mathcal{U}_P \times \mathbb{R} \times \mathbb{R}_+. \quad (59)$$

For each $z \in \mathcal{Z}$, let $\mathbb{P}_z^*$ denote the worst-case distribution induced by (57) (i.e., plugging $(U_P, \lambda_1, \lambda_2)$ into the density).

We define the VI operator $F : \mathcal{Z} \rightarrow \mathbb{R}^{m_P+2}$ as

$$F(z) \triangleq \begin{bmatrix} \nabla_{U_P} \mathbb{E}_{\mathbb{P}_z^*}[J_w(U_P, \mu)] \\ \mathbb{E}_{\mathbb{P}_z^*}[\mu] - \mu_0 \\ \mu_0^2 + \sigma^2 - \mathbb{E}_{\mathbb{P}_z^*}[\mu^2] \end{bmatrix}. \quad (60)$$

Find $z^* \in \mathcal{Z}$ such that

$$\langle F(z^*), z - z^* \rangle \geq 0, \quad \forall z \in \mathcal{Z}. \quad (61)$$

Equivalently, $\text{VI}(\mathcal{Z}, F)$ can be written as a monotone inclusion

$$0 \in F(z^*) + N_{\mathcal{Z}}(z^*), \quad (62)$$

where $N_{\mathcal{Z}}(\cdot)$ denotes the normal cone of $\mathcal{Z}$. The third component in (60) is written as $\mu_0^2 + \sigma^2 - \mathbb{E}[\mu^2]$ to correctly encode the second-moment upper bound together with $\lambda_2 \in \mathbb{R}_+$ via the normal cone inclusion.

5.3. Accelerated ADMM Solver for $\text{VI}(\mathcal{Z}, F)$

We apply a splitting scheme to (62) by introducing an auxiliary variable $y \in \mathcal{Z}$ and enforcing consensus:

$$z = y. \quad (63)$$

This yields a structured inclusion system suitable for ADMM-type iterations, where F is handled via a resolvent step and $N_{\mathcal{Z}}$ is handled via projection.

Let u be a scaled dual variable associated with (63), and let $\rho > 0$ be the penalty parameter. The (nonlinear) resolvent of F is defined by

$$J_{\rho F}(s) \triangleq (I + \rho F)^{-1}(s), \quad (64)$$

i.e., $z = J_{\rho F}(s)$ solves $z + \rho F(z) = s$.

The projection onto $\mathcal{Z}$ is

$$\Pi_{\mathcal{Z}}(q) \triangleq \arg \min_{v \in \mathcal{Z}} \|v - q\|_2^2. \quad (65)$$

The baseline ADMM-type splitting for $\text{VI}(\mathcal{Z}, F)$ is:

$$z^{k+1} = J_{\rho F}(y^k - u^k), \quad (66)$$

$$y^{k+1} = \Pi_{\mathcal{Z}}(z^{k+1} + u^k), \quad (67)$$

$$u^{k+1} = u^k + z^{k+1} - y^{k+1}. \quad (68)$$

Because $\mathcal{Z} = \mathcal{U}_P \times \mathbb{R} \times \mathbb{R}_+$, the projection in (67) decomposes into simple components:

$$\Pi_{\mathcal{Z}} \left(\begin{bmatrix} q_U \\ q_{\lambda_1} \\ q_{\lambda_2} \end{bmatrix} \right) = \begin{bmatrix} \Pi_{\mathcal{U}_P}(q_U) \\ q_{\lambda_1} \\ \max\{0, q_{\lambda_2}\} \end{bmatrix}. \quad (69)$$

To improve convergence speed under tight window-level iteration budgets, we optionally incorporate (i) over-relaxation and (ii) Nesterov-type extrapolation on the consensus/dual variables. These enhancements can be disabled, in which case the method reduces to the baseline iteration (66)–(68).

Let $\theta \in (0, 2)$ denote an over-relaxation parameter and define

$$\hat{z}^{k+1} \triangleq \theta z^{k+1} + (1 - \theta)y^k. \quad (70)$$

For Nesterov extrapolation, define a momentum sequence $\{t_k\}$ and coefficients $\{\alpha_k\}$ by

$$t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2}, \quad \alpha_k = \frac{t_k - 1}{t_{k+1}}, \quad t_0 = 1. \quad (71)$$

The extrapolated variables are

$$\bar{y}^k = y^k + \alpha_k(y^k - y^{k-1}), \quad \bar{u}^k = u^k + \alpha_k(u^k - u^{k-1}). \quad (72)$$

The accelerated ADMM iteration is then:

$$z^{k+1} = J_{\rho F}(\bar{y}^k - \bar{u}^k), \quad (73)$$

$$y^{k+1} = \Pi_{\mathcal{Z}}(\theta z^{k+1} + (1 - \theta)\bar{y}^k + \bar{u}^k), \quad (74)$$

$$u^{k+1} = \bar{u}^k + \theta z^{k+1} + (1 - \theta)\bar{y}^k - y^{k+1}. \quad (75)$$

The dominant cost lies in (73), which requires solving

$$z^{k+1} + \rho F(z^{k+1}) = \bar{y}^k - \bar{u}^k. \quad (76)$$

In practice, we exploit the block structure of F in (60). Writing $z = (U_P, \lambda_1, \lambda_2)$, (76) decomposes into:

$$U_P^{k+1} + \rho \nabla_{U_P} \mathbb{E}_{\mathbb{P}_{z^{k+1}}^*} [J_w(U_P^{k+1}, \mu)] = (\bar{y}^k - \bar{u}^k)_{U_P}, \quad (77)$$

$$\lambda_1^{k+1} + \rho (\mathbb{E}_{\mathbb{P}_{z^{k+1}}^*} [\mu] - \mu_0) = (\bar{y}^k - \bar{u}^k)_{\lambda_1}, \quad (78)$$

$$\lambda_2^{k+1} + \rho (\mu_0^2 + \sigma^2 - \mathbb{E}_{\mathbb{P}_{z^{k+1}}^*} [\mu^2]) = (\bar{y}^k - \bar{u}^k)_{\lambda_2}. \quad (79)$$

The expectations under $\mathbb{P}_{z^{k+1}}^*$ are computed from (57) using either (i) numerical quadrature on $[0, \mu_{\max}]$ or (ii) a sample-average approximation with importance weights induced by the exponential term. The resulting nonlinear system can be solved using a small number of safeguarded Newton or quasi-Newton iterations, which remains feasible due to the low-dimensionality of (λ_1, λ_2) and the window-localized horizon.

Assume that $\text{VI}(\mathcal{Z}, F)$ admits at least one solution and that F is monotone on $\mathcal{Z}$ and Lipschitz continuous with constant L_F . Under standard conditions for resolvent-based

splitting (with the resolvent $J_{\rho F}$ well-defined and single-valued), the baseline iteration (66)–(68) converges to a solution of $\text{VI}(\mathcal{Z}, F)$ for any $\rho \in (0, 2/L_F)$. If, in addition, F is strongly monotone on $\mathcal{Z}$ (which can be promoted by the entropic smoothing $\eta > 0$ together with convexity/regularity of J_w in U_P), then the baseline scheme admits a global linear (R-linear) convergence rate in standard metrics.

The accelerated variant (73)–(75) is used to improve practical speed under real-time constraints; when $\theta = 1$ and the extrapolation is disabled (i.e., $\alpha_k \equiv 0$), it reduces to the baseline method with the above convergence guarantee.

We employ primal and dual residuals to terminate iterations:

$$r^{k+1} \triangleq \|z^{k+1} - y^{k+1}\|_2, \quad s^{k+1} \triangleq \rho \|y^{k+1} - y^k\|_2, \quad (80)$$

and stop when both satisfy prescribed tolerances.

This section reformulates the window-wise VDRO optimality system into a finite-dimensional variational inequality $\text{VI}(\mathcal{Z}, F)$ by embedding the worst-case distribution through an entropy-regularized Gibbs mapping. An ADMM-type operator-splitting method is then developed, alternating between a resolvent step for the VI operator and a simple projection onto $\mathcal{Z}$, with optional over-relaxation and momentum to improve speed under window-level iteration budgets. In the next section, we provide numerical simulations to validate convergence behavior and quantify pursuit performance under varying interference regimes. Per-window solver implementation details are shown in Table 1, and the pseudo-code of Accelerated ADMM Operator-Splitting for $\text{VI}(\mathcal{Z}, F)$ is shown as Algorithm 2.

Table 1. Per-window solver implementation details.

Item	Typical Value	Notes
ADMM iterations per window	$N_{\text{it}} = 70$	fixed budget / early stop
Penalty parameter	$\rho = 0.9$	baseline
ρ sweep range	$[0.5, 1.3]$	sensitivity test
Stopping criterion	2×10^{-3}	primal residuals
Warm start	yes	from previous window

Algorithm 2 Accelerated ADMM Operator-Splitting for $\text{VI}(\mathcal{Z}, F)$ (window-wise)

Require: Window index w ; penalty $\rho > 0$; over-relaxation $\theta \in (0, 2)$; tolerances $\varepsilon_{\text{pri}}, \varepsilon_{\text{dual}}$; max iterations K_{max} ; initial (y^0, u^0) with $y^0 \in \mathcal{Z}$.

Ensure: Approximate solution z^* (and y^*).

```

1:  $y^{-1} \leftarrow y^0, \quad u^{-1} \leftarrow u^0, \quad t_0 \leftarrow 1.$ 
2: for  $k = 0, 1, \dots, K_{\text{max}} - 1$  do
3:    $t_{k+1} \leftarrow \frac{1 + \sqrt{1 + 4t_k^2}}{2}, \quad \alpha_k \leftarrow \frac{t_k - 1}{t_{k+1}}.$ 
4:    $\tilde{y}^k \leftarrow y^k + \alpha_k(y^k - y^{k-1}), \quad \tilde{u}^k \leftarrow u^k + \alpha_k(u^k - u^{k-1}).$   $\triangleright$  Resolvent step:  $z^{k+1} = J_{\rho F}(\tilde{y}^k - \tilde{u}^k)$ 
5:    $z^{k+1} \leftarrow J_{\rho F}(\tilde{y}^k - \tilde{u}^k).$   $\triangleright$  Over-relaxed consensus update
6:    $q^{k+1} \leftarrow \theta z^{k+1} + (1 - \theta)\tilde{y}^k + \tilde{u}^k.$ 
7:    $y^{k+1} \leftarrow \Pi_{\mathcal{Z}}(q^{k+1}).$   $\triangleright$  Scaled dual update
8:    $u^{k+1} \leftarrow \tilde{u}^k + \theta z^{k+1} + (1 - \theta)\tilde{y}^k - y^{k+1}.$   $\triangleright$  Stopping criteria
9:    $r^{k+1} \leftarrow \|z^{k+1} - y^{k+1}\|_2, \quad s^{k+1} \leftarrow \rho \|y^{k+1} - y^k\|_2.$ 
10:  if  $r^{k+1} \leq \varepsilon_{\text{pri}}$  and  $s^{k+1} \leq \varepsilon_{\text{dual}}$  then
11:    break
12:  end if
13: end for
14:  $z^* \leftarrow z^{k+1}, \quad y^* \leftarrow y^{k+1}.$ 

```

To improve readability, the end-to-end computational pipeline is summarized in Figure 4. The figure consolidates the data flow from online perception metrics (e.g., d_{JS} and the inferred interference intensity) to dynamic window generation, window-wise equilibrium construction, and the VI-based accelerated ADMM solver. This schematic provides a compact map of how the theoretical components developed in Sections 3–5 interact to yield a real-time implementable decision architecture.




Figure 4. Dynamic Time-Window Nash Equilibrium with VDRO and VI-ADMM.

6. Numerical Simulation

In this section, the simulations are organized to directly validate the motivating claims in Section 1: (i) dynamic windows improve responsiveness under rapidly varying perceptual degradation, (ii) VDRO hedges against ambiguity in interference intensity, and (iii) the proposed VI/ADMM solver achieves online feasibility via bounded per-window computation.

6.1. Simulation Setup and Evaluation Metrics

Numerical simulations are provided to validate the proposed VDRO-DTW framework in three-dimensional spacecraft pursuit–evasion engagements subject to time-varying, adversarial interference. The relative-motion dynamics, admissible control sets, and baseline cost structure (running and terminal terms) follow the modeling definitions in Section 2. Strategic incompleteness is introduced through the interference-corrupted belief update described in Section 2.3, and the perceptual degradation level is monitored online by the Jensen–Shannon (JS) divergence. The dynamic time-window mechanism in Section 3 is used to adapt the receding decision horizon, while the window-wise equilibrium is computed using the VDRO embedding and the VI/ADMM solver developed in Sections 4 and 5.

To ensure a fair and interpretable evaluation, the following metrics are reported consistently:

- Relative distance: $\|r(t)\|$, reflecting real-time pursuit effectiveness.
- Capture success: a trial is counted as successful if $\|r(t)\| \leq r_{\text{capt}}$ is achieved within the prescribed horizon.

- Pursuer control energy: $E_P \triangleq \int_0^T \|u_P(t)\|_2^2 dt$, reported as mean $\pm$ standard deviation for Monte Carlo studies.
- Minimum distance: $\min_{t \in [0, T]} \|r(t)\|$, capturing the closest approach even in partial-failure cases.

Unless stated otherwise, interference intervals are highlighted by gray shaded bands in trajectory plots; capture-zone spheres are displayed sparsely to avoid visual clutter.

For algorithmic comparison, three controllers are evaluated under identical dynamics, constraints, and interference realizations:

- FTW: a fixed time-window baseline with constant window length (no adaptation) and without distributional robustness.
- DTW: dynamic time-windowing (adaptive horizon) without the VDRO robustness layer.
- VDRO-DTW: the proposed method combining dynamic time windows with VDRO-based robustness in each window.

All methods are executed under the same numerical integration and window-boundary information flow; for fair timing and feasibility, the VI/ADMM solver is run with the same stopping criterion/iteration-budget policy whenever applicable (Section 5). Finally, the simulated relative separations and closing speeds are chosen to remain within the proximity-operation regime, so that the LVLH linearization underlying Equations (1)–(4) remains a valid approximation throughout the engagement.

To ensure full reproducibility, all numerical parameters and implementation settings used in the simulations are summarized in Tables 2–4, including orbital/model parameters, initial conditions, cost weights, window-trigger thresholds, VDRO ambiguity-set settings, and VI/ADMM solver hyperparameters.

Table 2. Dynamics and game-cost parameters used in simulations.

Item (symbol)	Value
T	200 s
Δt_{sim}	0.2 s
n	0.0011 s ^{−1}
$X_P(0)$	[0, 0, 0, 85.74, 85.74, 70] ^T
$X_E(0)$	[6123.7, 6123.7, 5000, 70.43, 70.43, 57.5] ^T
$U_P^{\text{max}}, U_E^{\text{max}}$	3.0, 2.0
α	1×10^{-4}
β_P, β_E	5×10^{-3} , 1.5×10^{-2}
$\eta(T)$	1×10^5
r_{capt}	2500 m

Table 3. Dynamic time-window and trigger parameters.

Item (symbol)	Value
$\Delta t_{\text{min}}, \Delta t_{\text{max}}$	0.8s, 2.0 s
Δt_{dwell}	0.6 s
$d_{\text{on}}, d_{\text{off}}$	0.08, 0.05
ρ_{switch}	0.12
δ_{state}	0.15
Δt_{hist}	20 s

Table 4. VDRO ambiguity-set and VI/ADMM solver settings.

Item (symbol)	Value
$\mu_{\max}$	0.8
η_{reg}	0.08
N_{alt}	3
N_{ADMM}	70
ρ	0.9
$\varepsilon_{\text{pri}}, \varepsilon_{\text{dual}}$	2×10^{-3} , 2×10^{-3}

6.2. Scenario-Based Verification Under Intermittent Interference

Scenario 1 (representative engagement with intermittent interference).

Figure 5 depicts a representative 3D engagement under intermittent interference using VDRO-DTW. Two qualitative behaviors are emphasized. First, geometric convergence is maintained: the pursuer trajectory steadily reduces the separation while remaining consistent with the capture-zone constraint. Second, interference-aware responsiveness is observed: during shaded interference intervals, the adaptive mechanism shortens the decision window, preventing long-horizon planning from becoming stale when perception is abruptly degraded.

The corresponding control inputs are shown in Figure 5b. The pursuer input evolves smoothly with bounded adjustments, while the evader input exhibits sharper variations during interference, consistent with an adversarial attempt to exploit degraded perception. The observed behavior supports the intended design principle: the pursuer reacts by tightening the decision cycle rather than over-amplifying instantaneous actuation.




Figure 5. Scenario 1: Three-dimensional pursuit–evasion engagement under interference (VDRO-DTW): (a) Relative trajectories with capture zone and interference intervals. (b) Control inputs of pursuer and evader.

Scenario 2 (more challenging geometry/stronger maneuver pressure).

A more challenging engagement configuration is reported in Figure 6. Compared with Scenario 1, the initial geometry and maneuver demands intensify the need for rapid policy correction and robustness to perception degradation. The resulting trajectory indicates that stable approach behavior is preserved without excessive detours, despite the stronger maneuver pressure. The control profiles further suggest that the pursuer increases actuation in a controlled manner to counteract bursts in the evader maneuver, while avoiding unrealistic discontinuities.




Figure 6. Scenario 2: Three-dimensional pursuit–evasion engagement under interference (VDRO-DTW): (a) Relative trajectories with capture zone and interference intervals. (b) Control inputs of pursuer and evader.

6.3. Monte Carlo Comparison Under Randomized Interference

To quantify robustness beyond individual realizations, Monte Carlo trials are performed under randomized interference sequences. The three strategies (FTW, DTW, and VDRO-DTW) are evaluated using identical trial sets to ensure paired comparability. Figure 7 reports the mean relative distance trajectories with one-standard-deviation bands over $N = 35$ trials. A rapid distance-reduction stage followed by a stabilization regime is observed across methods, while the dispersion bands reveal sensitivity to interference realizations.




Figure 7. Monte Carlo evaluation: mean relative distance with one-standard-deviation bands ($N = 35$).

Figure 8 summarizes capture success rate, pursuer energy, and minimum-distance statistics. Two points are highlighted. First, high capture success across methods indicates that time-windowing already provides strong responsiveness in the tested regime; DTW can be slightly more aggressive due to the absence of a robustness layer, whereas VDRO-DTW trades a small amount of nominal aggressiveness for improved protection against interference ambiguity. Second, the energy and minimum-distance statistics indicate that the proposed VDRO layer does not require disproportionate control effort to achieve robustness; the energy level remains comparable to the baselines, while the closest-approach

behavior is preserved under randomized disturbances. Overall, the Monte Carlo study supports that VDRO-DTW achieves a favorable robustness–efficiency balance without sacrificing core pursuit effectiveness.




Figure 8. Monte Carlo evaluation: capture success rate, control energy, and minimum-distance statistics (mean $\pm$ std).

6.4. Multi-Interference Stress Test and Resource Usage

Beyond single-pattern randomized interference, a mixed multi-interference setting is considered, where disturbance types and intensities switch over time. Figure 9 reports the relative distance evolution and cumulative pursuer energy. The distance remains bounded throughout the mixed-interference sequence, while the cumulative energy grows in a controlled manner. This result is consistent with the intended mechanism: the dynamic time-windowing maintains responsiveness under abrupt changes, and the VDRO layer mitigates performance degradation due to distributional ambiguity, jointly yielding a stable robustness–efficiency trade-off.




Figure 9. Multi-interference scenario: performance and resource usage: (a) Relative distance evolution under mixed interference. (b) Cumulative control energy consumption.

6.5. Dynamic Time-Window Adaptation Behavior

The time-varying window length t_w (equivalently, Δt_w) is visualized in Figure 10. The window consistently shrinks during high-degradation phases to enable faster policy refresh and expands during calmer phases to reduce unnecessary recomputation. This observation empirically confirms that the online window adaptation operates as designed: responsive-

ness is increased exactly when perception quality deteriorates, while computational load is reduced when stable information permits longer look-ahead.




Figure 10. Dynamic time-window adaptation in the presence of interference.

6.6. Rapidly Varying Interference Benchmark

We construct an interference profile with frequent on/off transitions and high-amplitude fluctuations in $d_{JS}(t)$, representing rapidly changing sensing denial. We compare the proposed dynamic-window policy against fixed-window baselines with identical per-window solvers and equal computation budgets. Performance is reported in terms of capture success rate, average interception time, and average computation time per window.

Figure 11 evaluates a rapidly varying sensing-denial benchmark with frequent on/off transitions and high-amplitude fluctuations. As shown in Figure 11a, the constructed profile induces recurrent bursts in $\mu(t)$ together with large oscillations in $d_{JS}(t)$, representing strongly non-stationary perceptual degradation. Under identical per-window solver settings and equal computation budgets, VDRO-DTW achieves the highest capture success rate in Figure 11b, indicating that the distributionally robust layer effectively mitigates decision degradation caused by short-lived uncertainty spikes. Moreover, conditional on capture, VDRO-DTW yields a lower interception-time distribution in Figure 11c, suggesting that robustness-aware planning not only improves feasibility but also accelerates convergence to capture when capture is achievable. Finally, Figure 11d shows that the computation time per window remains comparable across FTW, DTW, and VDRO-DTW, confirming that the performance gains are obtained without increased computational burden.

6.7. Sensitivity Analysis

We evaluate robustness of the conclusions to key design parameters. Specifically, we sweep (i) terminal-cost parameters (η, r_{capt}), (ii) attenuation-profile parameters (K_1, K_2, K_3) in (12), and (iii) solver penalty ρ . For each sweep, we report capture success rate, average interception time, and average computation time per window, showing that the proposed framework maintains consistent advantages over fixed-window baselines across a broad parameter range.

Figure 12 studies the robustness of the proposed framework with respect to key model and solver parameters. Figure 12a shows that the capture success rate varies smoothly as the attenuation-profile parameters (K_1, K_2, K_3) are jointly scaled, with only a mild performance decrease at more aggressive scaling, indicating that the proposed policy is not overly sensitive to moderate mismatch in the interference attenuation model. Figure 12b further demonstrates numerical robustness with respect to the ADMM penalty ρ : the success rate remains nearly constant over a broad range, while the computation time

per window exhibits a shallow minimum around intermediate ρ , consistent with typical ADMM behavior. Overall, the sensitivity results confirm that the observed performance advantages are maintained across a wide parameter range and do not rely on narrow hyperparameter tuning.




Figure 11. Rapidly varying interference benchmark under identical per-window computation budgets. (a) Rapid on/off interference profile visualized by $d_{JS}(t)$ and $\mu(t)$; (b) capture success rate; (c) interception time conditioned on successful capture; (d) computation time per decision window.




Figure 12. Sensitivity analysis of the proposed policy. (a) Capture success rate versus the scale factor applied to (K_1, K_2, K_3) in the attenuation profile; (b) sensitivity to the ADMM penalty ρ , reporting capture success rate (left axis) and computation time per window (right axis).

6.8. Perceptual Degradation Indicator and Inferred Interference Intensity

Finally, Figure 13 reports the perceptual degradation indicator (JS divergence) together with the inferred interference intensity. Their co-evolution provides an interpretable explanation of when and why the controller tightens the window and increases robustness: higher divergence indicates a stronger mismatch between nominal and interference-corrupted belief states, prompting more conservative, robustness-aware updates and shorter decision windows. This interpretability is important for contested-space autonomy, where performance must be explainable in addition to being effective.




Figure 13. Perceptual degradation and interference intensity: JS divergence and inferred interference level.

Across the representative scenarios, the Monte Carlo evaluation, and the mixed-interference stress tests, the results show that the proposed method behaves as intended. The pursuer maintains stable three-dimensional pursuit trajectories and adjusts its decision horizon when interference degrades perception, which is consistent with the goal of the dynamic time-window design. Compared with the fixed-window baseline and the dynamic-window baseline without VDRO, the proposed method preserves competitive capture behavior while keeping the pursuer control energy at a similar level. This suggests that the gain in robustness is not achieved by simply increasing actuation. In addition, the window-length adjustment and the divergence indicator evolve in a consistent way, which helps explain when the horizon is shortened and when robustness becomes more important. Overall, these results indicate that the VDRO-DTW framework reduces strategy staleness and addresses uncertainty in interference intensity, while remaining suitable for window-wise real-time operation.

Section 6 validates the proposed framework under representative engagement scenarios and interference regimes. Across all cases, the dynamic time-window policy yields improved outcomes over fixed-window baselines under comparable actuation limits and solver budgets, with the most noticeable gains occurring under rapidly varying perceptual degradation. This trend is consistent with the switching mechanism in Figure 2, where window boundaries are generated online from belief/state signals. Moreover, the VDRO component mitigates sensitivity to interference ambiguity and stabilizes performance across trials without introducing excessive conservatism. Finally, the reported iteration counts and runtimes confirm that the window-wise VI/ADMM implementation remains numerically stable under bounded per-window computation. Geometric visualizations in Figure 4 provide intuition for how window switching and allocation decisions evolve with degradation level.

7. Conclusions

This work studies spacecraft pursuit-evasion games in contested environments with time-varying interference and incomplete information and proposes a unified dynamic time-window (DTW) and variational distributionally robust optimization (VDRO) framework to mitigate strategy staleness and distributional ambiguity in the interference process. The method employs an evidence-based belief model and a Jensen–Shannon divergence indicator to quantify perceptual degradation online and to adapt the planning window accordingly—shortening the horizon during high-interference periods for rapid policy refresh and extending it in calmer phases to reduce unnecessary computation—while

a moment-constrained VDRO layer provides robustness against non-stationary, poorly characterized disturbances without excessive conservatism. Algorithmically, the VDRO KKT conditions are reformulated as a variational inequality and solved efficiently via an accelerated ADMM-type operator-splitting scheme suitable for windowed, real-time implementation. Numerical simulations and Monte Carlo studies demonstrate that the proposed approach yields stable three-dimensional engagement trajectories, competitive capture performance, and a favorable robustness-energy trade-off relative to fixed-window and non-robust baselines, while the window adaptation and divergence signals offer interpretable links between interference severity and control behavior. Future work will extend the framework to multi-agent engagements, higher-fidelity orbital and actuator constraints, and tighter integration of perception, interference identification, and learning-augmented robustness for onboard deployment.

Author Contributions: All authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by L.S., Z.H., Y.W., B.T. and P.H. The first draft of the manuscript was written by L.S. and all authors commented on previous versions of the manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Defense Science, Technology and Industry Bureau Stability Support Project (Grant HTKJ2024KL502029); the National Natural Science Foundation of Jiangsu Province of China (Grant BK20241393); the Aviation Science Foundation of China (Grant 2025L015052004); the General Program of National Natural Science Foundation of China (Grant 62373187).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: Author Binpeng Tian was employed by the company Aviation Industry Corporation of China, Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Ye, D.; Shi, M.; Sun, Z. Satellite proximate pursuit-evasion game with different thrust configurations. *Aerosp. Sci. Technol.* **2020**, *99*, 105715.
2. Li, J.; Li, C.; Zhang, Y. Guidance strategy of motion camouflage for spacecraft pursuit-evasion game. *Chin. J. Aeronaut.* **2024**, *37*, 312–319.
3. Han, C.; Chen, H.; Alonso, G.; Rao, Y.; Cubas, J.; Yin, J.; Wang, X. A linear model for relative motion in an elliptical orbit based on a spherical coordinate system. *Acta Astronaut.* **2019**, *157*, 465–476.
4. Drozd, K.; Furfaro, R.; Schiassi, E.; Johnston, H.; Mortari, D. Energy-optimal trajectory problems in relative motion solved via Theory of Functional Connections. *Acta Astronaut.* **2021**, *182*, 361–382.
5. Ostrowski, R.; Cywiriski, A.; Strzelec, M. Electronic warfare in the optical band: Main features, examples and selected measurement data. *Def. Technol.* **2021**, *17*, 1636–1649.
6. Sun, J.; Yuan, Y.; Greco, M. S.; Gini, F.; Yang, X. B.; Yi, W. Anti-deception jamming resource scheduling for multitarget tracking in distributed radar networks. *IEEE Trans. Aerosp. Electron. Syst.* **2024**, *60*, 9111–9128.
7. Song, H.; Xiao, M.; Xiao, J.; Liang, Y.; Yang, Z. A POMDP approach for scheduling the usage of airborne electronic countermeasures in air operations. *Aerosp. Sci. Technol.* **2016**, *48*, 86–93.
8. Bechini, M.; Gu, G.; Lunghi, P.; Lavagna, M. Robust spacecraft relative pose estimation via CNN-aided line segments detection in monocular images. *Acta Astronaut.* **2024**, *215*, 20–43.
9. Fu, S.; Gong, S.; Shi, P. Analytical pursuit-evasion game strategy in arbitrary Keplerian reference orbits. *Aerosp. Sci. Technol.* **2025**, *158*, 109946.
10. Fu, S.; Gong, S.; Wu, D.; Shi, P. Analytical solutions and winning conditions of elliptic-orbit target-attacker-defender game. *J. Guid. Control. Dyn.* **2026**, *49*, 257–268.
11. Pan, H.; Luo, M.; Wang, J.; Huang, T.; Sun, W. A Safe Motion Planning and Reliable Control Framework for Autonomous Vehicles. *IEEE Trans. Intell. Veh.* **2024**, *9*, 4780–4793.

12. Ohtsuka, T. A continuation/GMRES method for fast computation of nonlinear receding horizon control. *Automatica* **2004**, *40*, 563–574.
13. Zheng, J.; Sznaiier, M. Robust data-driven receding-horizon control for LQR with input constraints. *IEEE Control Syst. Lett.* **2025**, *9*, 2765–2770.
14. Wiesemann, W.; Kuhn, D.; Rustem, B. Robust markov decision processes. *Math. Oper. Res.* **2013**, *38*, 153–183.
15. Chamie, M. E.; Janak, D.; Aclkmese, B. Markov decision processes with sequential sensor measurements. *Automatica* **2019**, *103*, 450–460.
16. Hu, H.; Isele, D.; Bae, S.; Fisac, J. F. Active uncertainty reduction for safe and efficient interaction planning: A shielding-aware dual control approach. *Int. J. Robot. Res.* **2023**, *43*, 1382–1408.
17. Zhou, P.; Li, S.; Zhao, B.; Wahlberg, B.; Hu, X. Nature-inspired dynamic control for pursuit-evasion of robots. *Automatica* **2026**, *183*, 112629.
18. Chen, G.; Xu, W.; Li, Z.; Liu, Y.; Liu, X. Research on the multi-robot cooperative pursuit strategy based on the zero-sum game and surrounding points adjustment. *Machines* **2021**, *9*, 187.
19. Zhao, Y.; Guo, C.; Mi, J.; Wang, L.; Wang, H.; Zhang, H. Research on path tracking technology for tracked unmanned vehicles based on DDPG-PP. *Machines* **2025**, *13*, 603.
20. Su, Z.; Zheng, S.; Xu, Z.; Cheng, L.; Tao, C.; Qie, R.; Feng, W.; Zhang, Z.; Xu, Y. Evade unknown pursuer via pursuit strategy identification and model reference policy adaptation (MRPA) algorithm. *Drones* **2024**, *8*, 655.
21. Delage, E.; Ye, Y. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* **2010**, *58*, 595–612.
22. Goh, J.; Sim, M. Distributionally robust optimization and its tractable approximations. *Oper. Res.* **2010**, *58*, 902–917.
23. Mohajerin Esfahani, P.; Kuhn, D. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable Reformulations. *Math. Program.* **2018**, *171*, 115–166.
24. Dempster, A. P. Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Stat.* **1967**, *38*, 325–339.
25. Lin, J. Divergence measures based on the shannon entropy. *IEEE Trans. Inf. Theory* **1991**, *37*, 145–151.
26. Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **2011**, *3*, 1–122.
27. Touma, T.; Daş, E.; Tevere, E.; Feather, M.; Kolcio, K.; Prather, M.; Candela, A.; Goel, A.; Kramer, E.; Nayar, H. AI space cortex: An experimental system for future era space exploration. *IEEE Trans. Field Robot.* **2025**, *2*, 601–625.
28. Carrasco, A.; Rodriguez-Fernandez, V.; Linares, R. Large language models as autonomous spacecraft operators in kerbal space program. *Adv. Space Res.* **2025**, *76*, 3480–3497.
29. Xia, R.; Xiao, F. A complex plausibility belief Jensen-Shannon divergence with its application in multi-source information fusion. *Eng. Appl. Artif. Intell.* **2024**, *137*, 109056.

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