Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Dynamic Belief State Planning Framework for Ambush Avoidance in Contested Environments: A Game-Theoretic Approach

Ivan Enzo Gargano

igargano@unibz.it

Free University of Bozen-Bolzano

Karl Dietrich von Ellenrieder

Free University of Bozen-Bolzano

Marianna Vivolo

Iveco Defence Vehicles

Research Article

Keywords: adversarial path planning, dynamic belief state planning, game-theoretic planning under uncertainty, temporal belief dynamics, pursuit-evasion games, information decay modeling

Posted Date: February 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8916440/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Dynamic Belief State Planning Framework for Ambush Avoidance in Contested Environments: A Game-Theoretic Approach

Ivan Enzo Gargano[1,2], Karl D. von Ellenerieder[1], Marianna Vivolo[2]

[1]Facoltà di Ingegneria, Libera Università di Bolzano, Piazza Università, 1, Bolzano, 39100, Italy.
[2]Iveco Defence Vehicles, Via Alessandro Volta, 6, Bolzano, 39100, Italy.

Contributing authors: IvanEnzo.Gargano@student.unibz.it; karl.vonellenrieder@ieee.org; marianna.vivolo@ivecogroup.com;

Abstract

Existing approaches to adversarial path planning for ambush avoidance assume perfect, static information about opponent strategies and environmental conditions—an assumption that fails in dynamic settings where information degrades and adaptive agents learn from observed behavioral patterns. We present a dynamic belief state planning framework that extends game-theoretic path planning through four multiplicative adjustment factors modeling temporal information decay, opponent adaptation through strategy learning, environmental uncertainty propagation, and partial observability constraints. The framework operates in belief space rather than state space while maintaining polynomial-time computational complexity through linear programming formulations with temporally-adjusted utility values. Validation across 180 simulated scenarios demonstrates that our dynamic belief planning approach maintains 95% of optimal performance under severe information degradation conditions where static game-theoretic approaches degrade to 6% effectiveness—representing an 89 percentage point improvement in decision quality. The framework establishes critical operational thresholds for autonomous agents: active information gathering becomes beneficial above environmental uncertainty U=0.36, behavioral strategies with entropy below 0.3 enable opponent exploitation through learning, and extended planning horizons exceeding 20 decision epochs require 15-20% performance degradation margins. The framework provides computationally tractable

solutions for autonomous agent deployment while bridging theoretical optimality and practical realizability for agents operating under degraded observability and adaptive opponents.

Keywords: adversarial path planning, dynamic belief state planning, game-theoretic planning under uncertainty, temporal belief dynamics, pursuit-evasion games, information decay modeling

1 Introduction

Autonomous navigation in adversarial multi-agent environments represents a fundamental challenge in agent decision-making where movement becomes a strategic interaction between intelligent entities. Unlike agents operating in cooperative or neutral environments that optimize for efficiency, autonomous agents in adversarial settings must achieve their objectives while countering opponents who observe, learn from, and exploit predictable behavioral patterns [1]. This challenge manifests across diverse application domains including logistics routing, resource transportation, and autonomous vehicle coordination, where behavioral predictability creates strategic vulnerabilities that adversarial agents can exploit [2].

The dominant approach to addressing adversarial navigation is *Ambush Avoidance*, which formulates route planning as a multi-agent optimization problem where autonomous agents must traverse networks while minimizing exposure to adversarial interdiction attempts [3]. Ambush avoidance has emerged as a primary paradigm because it directly captures the strategic nature of interactions between autonomous agents and intelligent adversaries. Contemporary Ambush Avoidance approaches rely on game-theoretic frameworks modeling strategic interactions as *two-player zero-sum games*, achieving significant theoretical advances through linear programming solutions with proven computational efficiency [4, 5]. However, these approaches operate under three critical idealizations that limit practical applicability in multi-agent systems: (1) perfect and static information assumptions about network topology and threat distributions; (2) treatment of adversarial agents as static entities rather than learning agents with adaptation capabilities; and (3) complete neglect of temporal dynamics where information loses value as environmental conditions evolve over time.

This paper addresses these limitations through a *dynamic belief state planning framework* (Enhanced Dynamic Belief, EDB) that operates in belief space rather than state space, explicitly modeling temporal uncertainty, adversary adaptation through learning, and information degradation. The key novelty of our approach lies in the integration of temporal dynamics into game-theoretic path planning through four multiplicative adjustment factors that independently model distinct information degradation processes—temporal decay, opponent learning, environmental uncertainty propagation, and partial observability—while preserving polynomial-time computational complexity. Unlike existing belief state planning methods that treat uncertainty as static, or reinforcement learning approaches that require extensive training data, our

framework provides tractable closed-form solutions that adapt continuously as information ages and adversaries learn, bridging the gap between theoretical game-theoretic optimality and practical deployment under degraded observability conditions.

Our approach integrates four key innovations: *Temporal Belief Decay Modeling* formalizes information reliability degradation over time; *Adversary Adaptation Framework* models opponent learning from observed agent behavioral patterns through strategy entropy quantification; *Dynamic Uncertainty Integration* combines multiple environmental and operational uncertainty sources through principled propagation; and *Information Staleness Quantification* establishes formal methods for evaluating multi-source information with varying reliability and age characteristics.

Our framework fundamentally shifts from traditional state space planning to belief space planning, where agent decisions account for uncertainty evolution over time while maintaining computational efficiency through multiplicative adjustment factors that modify base threat assessments within the linear programming structure. We validate our approach through comprehensive simulation studies across 180 experimental scenarios, demonstrating that our Enhanced Dynamic Belief approach maintains 95% of optimal performance under severe information degradation while static approaches degrade to 6% performance—representing an 89 percentage point improvement in decision quality with 51% improvement in behavioral unpredictability that prevents opponent exploitation through learning.

The remainder of this paper is organized as follows. Section 2 reviews existing approaches to ambush avoidance and belief state planning in multi-agent systems. Section 3 presents the mathematical formulation of our dynamic belief state planning framework. Section 4 provides experimental validation demonstrating effectiveness compared to existing methods. Section 5 discusses implications and limitations for multi-agent systems, and Section 6 summarizes contributions and future research directions.

2 Related Work

Game-theoretic approaches provide foundational frameworks for strategic navigation in adversarial multi-agent environments, where autonomous agents must navigate while accounting for intelligent opponents. Classical formulations model these scenarios as *two-player zero-sum games*, achieving computational efficiency through linear programming solutions [4, 5]. Joseph and Feron [6] formulated route selection as a minimax optimization problem, demonstrating that optimal solutions require mixed strategies with complexity $\mathcal{O}(|E||N|)$. Karaman and Frazzoli [7] extended this framework with hierarchical Stackelberg games, while Salani et al. [8] addressed multi-destination scenarios through flow-based programming. Recent advances integrate reinforcement learning with game-theoretic foundations, including double oracle algorithm approaches [9] and Yao et al.'s [10] WoLF-BSS-Q algorithm for dynamic strategy adaptation.

These approaches establish rigorous mathematical foundations for agent decision-making under strategic interaction. However, they typically assume perfect, static

information about environment topology, opponent capabilities, and threat distributions [11]. While this assumption enables tractable solutions, it presents opportunities for extension to scenarios where information quality degrades over time and environmental conditions evolve dynamically—characteristics common in real-world multi-agent deployments where agents operate with limited and aging observations.

2.1 Belief State Planning Under Uncertainty

Recognizing that autonomous agents typically operate with incomplete and uncertain information, researchers have developed belief state planning approaches using Partially Observable Markov Decision Processes (POMDPs). These frameworks address agent navigation under uncertainty through methods such as the Feedback Information Roadmap (FIRM) [12–14]. Recent extensions integrate temporal logic constraints [15] and simplified belief guides [16] to enhance computational tractability. Multi-source information fusion has advanced through automated frameworks [17] and comprehensive integration methods [18].

While these approaches effectively handle spatial uncertainty and incomplete observability, they generally model uncertainty as static throughout agent operation. This presents an opportunity to explore how belief states evolve temporally—specifically, how information value naturally decreases over time and how uncertainty grows in the absence of fresh observations [19]. Incorporating temporal dynamics into belief state representations could enhance agent decision-making in extended-duration tasks where initial information ages.

2.2 Control-Theoretic Methods for Agent Dynamics

Control-theoretic approaches offer rigorous mathematical foundations through Hamilton-Jacobi-Isaacs (HJI) formulations for multi-agent interaction scenarios. Building on Isaacs' seminal work [20], Weintraub et al. [21] provided comprehensive analysis of pursuit-evasion differential games with detailed problem classifications.

Neural network approaches address computational challenges in high-dimensional state spaces. Nakamura-Zimmerer et al. [22] introduced data-driven methods for solving high-dimensional Hamilton-Jacobi-Bellman equations without discretization, while Fisac et al. [23] combined HJ reachability with reinforcement learning for safety-critical agent systems. Continuous environment navigation has benefited from visual modeling approaches, with Meng et al. [24] developing TerrainNet for complex terrain understanding.

These methods provide powerful frameworks for continuous-time agent dynamics. An opportunity exists to extend these approaches to incorporate dynamic environmental changes and discrete information updates that characterize many multi-agent scenarios, where agents receive periodic observations from multiple sources with varying reliability characteristics.

2.3 Dynamic and Adaptive Multi-Agent Approaches

Recent work addresses dynamic elements in adversarial multi-agent scenarios. Masters and Sardina [25] established metrics for deceptive agent behavior including Last

Deceptive Point and Radius of Maximum Probability. Machine learning integration has advanced through graph neural network policies [26] and Value of Information objectives [27], enabling agents to reason about new information gathering value.

Recent security-game formulations model agent routing under adversarial interdiction through randomized equilibria. Contested-routing approaches [3, 28] cast agent-opponent interactions as two-player zero-sum games, computing randomized equilibrium strategies with scalable double-oracle best responses. These works advance adversarial modeling while maintaining computational tractability.

These approaches enhance opponent modeling and strategic reasoning. However, they typically operate with per-instance static information, presenting opportunities to incorporate temporal information degradation and multi-source intelligence staleness into the agent's decision-making process—aspects we address in this work.

2.4 Research Motivation and Contributions

Our analysis of existing literature reveals several promising research directions that motivate our framework. We identify opportunities to extend current approaches in four key areas:

- **Temporal evolution of uncertainty**: Current approaches typically model uncertainty as static, presenting an opportunity to formalize how information reliability naturally degrades over time and how uncertainty accumulates in the absence of new observations. This temporal dimension is particularly relevant for long-duration agent operations where initial information ages and environmental conditions evolve.
- **Opponent adaptation modeling**: Most existing frameworks treat opponents as stationary entities with fixed strategies. However, in multi-mission scenarios, intelligent opponents may learn from observed agent behavior patterns. Incorporating opponent learning dynamics could enhance agent decision-making by accounting for strategic adaptation over repeated interactions.
- **Integrated uncertainty treatment**: While existing methods address environmental uncertainty, observational limitations, and opponent modeling separately, there is potential value in treating these uncertainty sources in an integrated manner that captures their temporal evolution and interaction in dynamic multi-agent environments.
- **Information staleness quantification**: Multi-source information fusion research acknowledges data quality variations, yet formal metrics for quantifying information age and reliability decay could be more tightly integrated into agent planning algorithms. This would enable more principled decision-making when agents must act based on observations of varying age and quality.

These research directions collectively motivate a unified framework that addresses temporal uncertainty evolution, opponent learning dynamics, and multi-source information integration within a coherent mathematical structure. Our work aims to preserve computational tractability while enhancing the realism of agent models for deployment in dynamic, adversarial multi-agent environments where information degrades over time and opponents adapt to observed behaviors.

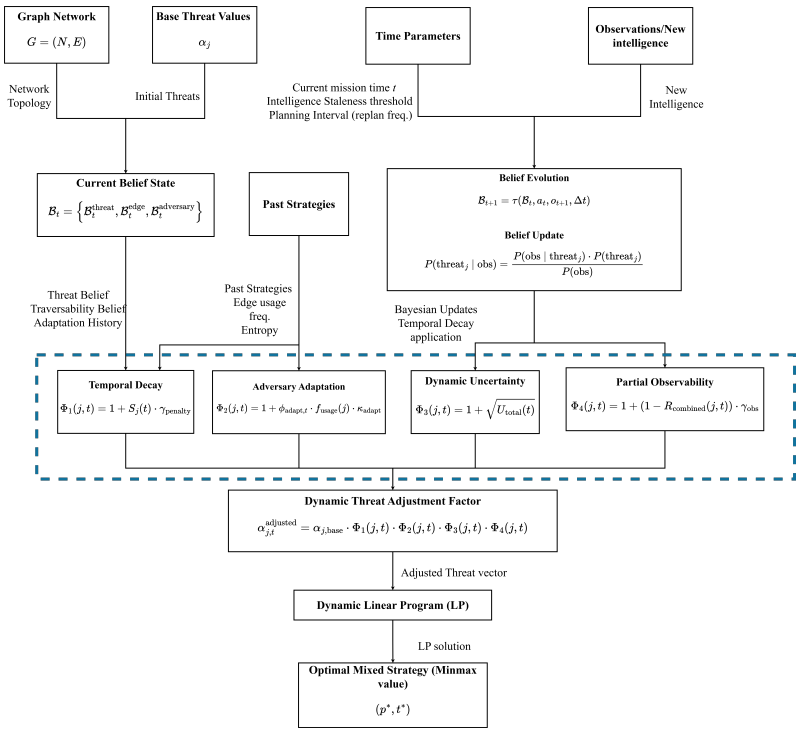


Fig. 1 Enhanced Dynamic Belief (EDB) framework architecture for autonomous agent decision-making. The system processes network topology, base environmental costs α_j, temporal parameters, and agent observations through belief state management, producing four temporal adjustment factors: temporal decay Φ_1 (information staleness), adversary adaptation Φ_2 (opponent learning from observed strategies), dynamic uncertainty Φ_3 (environmental variability), and partial observability Φ_4 (sensor limitations). These factors multiplicatively adjust cost values $\alpha_{j,t}^{\text{adjusted}}$ before feeding into a dynamic optimization procedure that computes optimal mixed strategies. Unlike static approaches, EDB continuously updates environmental assessments through independent information streams modeling distinct processes in multi-agent adversarial scenarios, maintaining polynomial-time complexity suitable for real-time agent deployment.

3 Methodology

We extend classical game-theoretic path planning to address dynamic uncertainty through belief space planning with temporal evolution. Our framework transforms static environmental assessments into time-varying beliefs that decay, adapt, and compound through four multiplicative adjustment factors, enabling autonomous agents to reason effectively in adversarial multi-agent environments, as illustrated in Figure 1.

3.1 Problem Formulation

3.1.1 Static Game-Theoretic Foundation.

The adversarial navigation problem models path selection as a two-player zero-sum game [29]. A navigating agent traverses network $G = (N, E)$ from origin A to destination B, while an adversarial agent attempts to intercept along the path. Each node $n_j \in N$ carries an interception cost value α_j representing expected utility loss if the adversary successfully interdicts at that location.

Both agents employ mixed strategies [4, 5]: the navigating agent's strategy p assigns probability p_{ij} to edge e_{ij}, while the adversarial agent's strategy q assigns probability q_j to positioning at node n_j. The solution concept is the Nash Equilibrium (NE) [30, 31], where neither agent can unilaterally improve their expected payoff. The expected outcome for the adversarial agent is:

$$V = \sum_{j \in N} q_j \cdot \alpha_j \cdot \sum_{(i,j) \in E} p_{ij}. \tag{1}$$

As the optimization objective, the navigating agent seeks to minimize the maximum expected cost:

$$p^* = \arg\min_{p} \left(\max_{j \in N} \sum_{i:(i,j) \in E} p_{ij} \cdot \alpha_j \right). \tag{2}$$

This Minimax problem [32] can be formulated as a linear program (LP):

$$\min_{p,z} \quad z \tag{3}$$

$$\text{subject to} \quad \sum_{i:(i,j) \in E} p_{ij} \cdot \alpha_j \leq z, \qquad \forall j \in N \tag{4}$$

$$\sum_{i:(i,j) \in E} p_{ij} = \sum_{k:(j,k) \in E} p_{jk}, \qquad \forall j \in N \setminus \{A, B\} \tag{5}$$

$$\sum_{j:(A,j) \in E} p_{Aj} = 1 \tag{6}$$

$$\sum_{i:(i,B) \in E} p_{iB} = 1 \tag{7}$$

$$p_{ij} \geq 0. \qquad \forall (i,j) \in E \tag{8}$$

3.1.2 Dynamic Extension.

Static formulations assume perfect, unchanging environmental knowledge—unrealistic in multi-agent adversarial scenarios where information degrades and opponents adapt.

We replace static α_j with time-varying $\alpha_{j,t}^{\text{adjusted}}$:

$$\alpha_{j,t}^{\text{adjusted}} = \alpha_{j,\text{base}} \cdot \prod_{k=1}^{4} \Phi_k(j,t), \tag{9}$$

where adjustment factors Φ_k with $k \in [1,4]$ capture temporal information decay, opponent learning, uncertainty growth, and observability limitations.

3.2 Belief State Framework

Operating in belief space enables agent decision-making under evolving uncertainty [33, 34]. The belief state at time t encompasses all uncertain information available to the agent:

$$\mathcal{B}_t = \left\{ \mathcal{B}_t^{\text{threat}}, \mathcal{B}_t^{\text{edge}}, \mathcal{B}_t^{\text{adversary}} \right\} \tag{10}$$

where $\mathcal{B}_t^{\text{threat}}$ contains probability distributions over adversarial positions used in computing Φ_1 and Φ_4, $\mathcal{B}_t^{\text{edge}}$ tracks environmental traversability conditions affecting Φ_3, and $\mathcal{B}_t^{\text{adversary}}$ maintains opponent adaptation history for Φ_2.

The belief state evolves through temporal dynamics and agent observations [35]:

$$\mathcal{B}_{t+1} = \tau(\mathcal{B}_t, a_t, o_{t+1}, \Delta t) \tag{11}$$

where τ is the belief update function incorporating agent actions a_t, observations o_{t+1}, and elapsed time Δt. This evolution directly drives the adjustment factors Φ_k with $k \in [1,4]$ that transform static cost values into dynamic assessments, as detailed in the following subsections.

3.3 Temporal Adjustment Factors

3.3.1 Temporal Decay.

Information quality degrades exponentially over time [36]

$$\mathrm{r}(t) = r_0 \cdot e^{\lambda(t - t_{\text{obs}})} \tag{12}$$

where $r_0 \in [0,1]$ is the initial information quality at observation time t_{obs} and $\lambda \in (0,1)$ represents the information decay rate. Information staleness (how outdated the current information is) is quantified as

$$S_j(t) = \frac{t - t_{\text{last_obs}}(j)}{T_{\max}} \tag{13}$$

normalized by maximum useful information lifetime $T_{\max}$ [37]. As information ages, uncertainty grows:

$$\sigma_t^2 = \sigma_{\text{obs}}^2 + \gamma \cdot (t - t_{\text{obs}}), \tag{14}$$

where γ represents the rate of uncertainty growth. We model threat beliefs using Beta distributions, which provide natural conjugate priors for binary outcomes [38]. As information ages without new observations, uncertainty grows through Beta parameter drift:

$$a_{t+1} = a_t + \delta_{\text{uncertainty}} \cdot \Delta t \tag{15}$$
$$b_{t+1} = b_t + \delta_{\text{uncertainty}} \cdot \Delta t, \tag{16}$$

where (a_t, b_t) are the Beta distribution parameters at time t. The temporal adjustment factor Φ_1 becomes:

$$\Phi_1(j,t) = 1 + S_j(t) \cdot \gamma_{\text{penalty.}} \tag{17}$$

3.3.2 Adversary Adaptation.

Opponent agents learn from observing behavioral patterns, adapting their strategies based on historical observations [39]. Low-entropy strategies are particularly vulnerable to such exploitation, as their predictability enables opponent optimization. The adaptation belief evolves as

$$\phi_{\text{adapt},t+1} = \min(0.95, \phi_{\text{adapt},t} + \eta \cdot (1 - H(s_t)) \cdot \Delta t), \tag{18}$$

where $\phi_{\text{adapt},t}$ represents the current opponent adaptation level, η is the learning rate, and $H(s_t)$ measures the entropy of strategy s_t, quantifying its predictability. Adversarial positioning distributions adjust based on historical path usage:

$$\alpha_{\text{threat},j,t+1} = \alpha_{\text{threat},j,t} + \phi_{\text{adapt},t} \cdot f_{\text{usage}}(j) \cdot \beta_{\text{adjust}}, \tag{19}$$

where $f_{\text{usage}}(j)$ represents the historical frequency with which node j has been visited and β_{adjust} is the adjustment rate for adversarial concentration. The adaptation factor Φ_2 becomes:

$$\Phi_2(j,t) = 1 + \phi_{\text{adapt},t} \cdot f_{\text{usage}}(j) \cdot \kappa_{\text{adapt}}. \tag{20}$$

3.3.3 Dynamic Uncertainty.

Multiple uncertainty sources compound over time in multi-agent environments [37, 40]. Environmental uncertainty captures dynamics such as changing conditions and infrastructure degradation

$$U_{\text{env}}(t) = w_{\text{weather}} \cdot \sigma^2_{\text{weather}}(t) + w_{\text{deg}} \cdot \frac{t}{T_{\max}}, \tag{21}$$

where w_{weather} weights environmental variability impact, $\sigma^2_{\text{weather}}(t)$ is environmental variance, and w_{deg} weights degradation impact (default: 0.2). Adversary mobility

uncertainty grows with opponent movement capability:

$$U_{\text{mobility}}(t) = \theta_{\text{mobility}} \cdot \Delta t, \tag{22}$$

where θ_{mobility} represents opponent mobility rate. Observation uncertainty stems from sensor limitations:

$$U_{\text{obs}}(t) = \varepsilon_{\text{fp}} + \varepsilon_{\text{fn}}, \tag{23}$$

where ε_{fp} and ε_{fn} are false positive and negative rates characterizing sensor performance [41]. Temporal uncertainty accumulates between observations

$$U_{\text{temporal}}(t) = \gamma \cdot (t - t_{\text{last_obs}}), \tag{24}$$

where γ is the uncertainty growth rate. These uncertainties also affect edge traversability, which degrades over time [42]:

$$\alpha_{\text{edge},t+1} = \alpha_{\text{edge},t} \cdot (1 - r_{\text{deg}} \cdot \Delta t - w_{\text{sensitivity}} \cdot \xi_t), \tag{25}$$

where r_{deg} is infrastructure degradation rate, $w_{\text{sensitivity}}$ captures environmental vulnerability, and $\xi_t \sim \mathcal{N}(0, \sigma^2_{\text{weather}})$ is environmental noise. The combined uncertainty adjustment factor becomes [43]:

$$\Phi_3(j, t) = 1 + \sqrt{U_{\text{total}}(t)}, \tag{26}$$

where total uncertainty aggregates all sources:

$$U_{\text{total}}(t) = U_{\text{env}}(t) + U_{\text{mobility}}(t) + U_{\text{obs}}(t) + U_{\text{temporal}}(t). \tag{27}$$

3.3.4 Partial Observability.

Autonomous agents rely on imperfect sensors with distance-dependent detection probability:

$$P_{\text{detect}}(d) = P_{\max} \cdot e^{-d/R_{\text{sensor}}}, \tag{28}$$

where d is target distance and R_{sensor} is effective sensor range [44, 45]. Multi-source information fusion weights reliability by both source quality and information age [35]:

$$R_{\text{combined}} = \frac{\sum_i w_i \cdot r_i \cdot (1 - S_i)}{\sum_i w_i} \tag{29}$$

where w_i weights source importance, r_i is base reliability (e.g., aerial sensor: 0.85, mobile sensor: 0.90, stationary sensor: 0.60), and $S_i = (t - t_{\text{obs},i})/T_{\max}$ quantifies staleness. When new observations arrive, beliefs update via Bayes' rule [46]:

$$P(\text{threat}_j \mid \text{obs}) = \frac{P(\text{obs} \mid \text{threat}_j) \cdot P(\text{threat}_j)}{P(\text{obs})}. \tag{30}$$

Between observations, beliefs revert toward prior distributions based on staleness:

$$p_{\text{threat}}(t) = p_{\text{current}} + \frac{t - t_{\text{obs}}}{T_{\max}} \cdot \gamma_{\text{staleness}} \cdot (p_{\text{prior}} - p_{\text{current}}), \tag{31}$$

where $\gamma_{\text{staleness}}$ controls reversion rate. The observability factor becomes:

$$\Phi_4(j,t) = 1 + (1 - R_{\text{combined}}(j,t)) \cdot \gamma_{\text{obs}} \tag{32}$$

3.4 Integrated Dynamic Linear Program

The dynamic formulation extends the static LP (Equations 3-8) by replacing static α_j with time-adjusted values that integrate all four temporal factors:

$$\min_{p,z} \quad z \tag{33}$$

$$\text{subject to} \quad \sum_{i:(i,j)\in E} p_{ij} \cdot \alpha_{j,t}^{\text{adjusted}} \leq z, \qquad \forall j \in N \tag{34}$$

$$\sum_{i:(i,j)\in E} p_{ij} = \sum_{k:(j,k)\in E} p_{jk}, \qquad \forall j \in N \setminus \{A, B\} \tag{35}$$

$$\sum_{j:(A,j)\in E} p_{Aj} = 1 \tag{36}$$

$$\sum_{i:(i,B)\in E} p_{iB} = 1 \tag{37}$$

$$p_{ij} \geq 0, \qquad \forall (i,j) \in E. \tag{38}$$

The key transformation lies in the adjusted cost values:

$$\alpha_{j,t}^{\text{adjusted}} = \alpha_{j,\text{base}} \cdot \Phi_1(j,t) \cdot \Phi_2(j,t) \cdot \Phi_3(j,t) \cdot \Phi_4(j,t) \tag{39}$$

where:

- **Temporal decay (Eq. 17):**

$$\Phi_1(j,t) = 1 + \frac{t - t_{\text{last_obs}}(j)}{T_{\max}} \cdot \gamma_{\text{penalty}}$$

  models information degradation over time (cf. Eqs. 12–16).
- **Opponent learning (Eq. 20):**

$$\Phi_2(j,t) = 1 + \phi_{\text{adapt},t} \cdot f_{\text{usage}}(j) \cdot \kappa_{\text{adapt}}$$

  captures adversarial agent adaptation based on strategy entropy (cf. Eqs. 18–19).

- **Uncertainty growth (Eq. 26):**

$$\Phi_3(j,t) = 1 + \sqrt{U_{\text{total}}(t)}$$

  represents compounding uncertainty from multiple sources in the multi-agent environment (cf. Eqs. 21–25, Eq. 27).
- **Observability limits (Eq. 32):**

$$\Phi_4(j,t) = 1 + (1 - R_{\text{combined}}(j,t)) \cdot \gamma_{\text{obs}}$$

  accounts for multi-source sensor fusion with varying reliability (cf. Eqs. 28–31).

An overview of the complete agent decision-making framework is provided in Algorithm 1.

4 Results

We validate our dynamic belief state planning framework through systematic simulation studies that quantify the value of temporal uncertainty modeling in adversarial multi-agent environments. All experiments were conducted on an Intel Core i7-12800H system with 14 cores, 20 threads, 24MB L3 cache, and 32GB RAM running Ubuntu 22.04.5 LTS, ensuring reproducible computational performance analysis.

4.1 Experimental Design Overview

Our validation isolates each algorithmic innovation through progressive comparison of four approaches representing the evolution from deterministic to uncertainty-aware planning in adversarial settings:

- **Dijkstra Baseline:** Pure distance optimization without adversarial consideration, establishing fundamental navigation performance bounds in the underlying graph structure.
- **Static Game-Theoretic (SGT):** Established two-player zero-sum formulation providing optimal mixed strategies under perfect information assumptions [4, 5].
- **Basic Belief State (BBS):** Game-theoretic optimization extended to probabilistic beliefs over uncertain conditions, treating uncertainty as static throughout agent execution.
- **Enhanced Dynamic Belief (EDB):** Our complete framework integrating temporal belief decay, adversary adaptation, uncertainty evolution, and information staleness quantification as formalized in Section 3, using the parameters defined in Table 1.

Parameter calibration reflects realistic agent constraints in adversarial environments. Threat mobility parameters with $\sigma \in \{0, 0.1, 0.3\}$ represent static positions, limited mobility ranges of 5 to 10 spatial units, and high mobility ranges of 30 to 50 spatial units [47, 48]. The information decay factor $\lambda = 0.95$ corresponds to 5% degradation per time unit observed in extended agent operations [49]. Network configurations ranging from 15 to 40 nodes span local to distributed environments with

Algorithm 1 Enhanced Dynamic Belief (EDB) Framework for Agent Decision-Making

Require: $G(N, E)$, $\alpha_{j,\text{base}}$, T_{mission}, $T_{\max}$, Δt
Ensure: (p^*, z^*)
$\mathcal{B}_0 \leftarrow$ InitializeBeliefs() ▷ Eq. 10
$\mathcal{H} \leftarrow \emptyset$; $t_{\text{obs}}(j) \leftarrow 0, \forall j$
for $t = 0$ **to** T_{mission} **step** Δt **do**
  // Belief Evolution
  $\mathcal{B}_t \leftarrow \tau(\mathcal{B}_{t-1}, a_t, o_{t+1}, \Delta t)$ ▷ Eq. 11
  if $o_t \neq \emptyset$ **then**
    Apply Bayesian update ▷ Eq. 30
    Update $t_{\text{obs}}(j)$ for observed nodes
  end if
  Apply temporal decay to Beta parameters ▷ Eqs. 15–16
  // Compute Adjustment Factors
  for each $j \in N$ **do**
    $S_j \leftarrow$ StalenessFunction(j, t) ▷ Eq. 13
    $\Phi_1 \leftarrow$ TemporalDecay(S_j) ▷ Eq. 17
    $f_{\text{usage}} \leftarrow$ UsageHistory$(\mathcal{H}, j)$
    $H \leftarrow$ StrategyEntropy$(\mathcal{H})$
    Update ϕ_{adapt} ▷ Eq. 18
    $\Phi_2 \leftarrow$ OpponentAdaptation$(\phi_{\text{adapt}}, f_{\text{usage}})$ ▷ Eq. 20
    $U_{\text{total}} \leftarrow$ CombinedUncertainty(t) ▷ Eq. 27
    $\Phi_3 \leftarrow$ UncertaintyFactor(U_{total}) ▷ Eq. 26
    $R_{\text{combined}} \leftarrow$ ReliabilityFusion(j, t) ▷ Eq. 29
    $\Phi_4 \leftarrow$ ObservabilityFactor(R_{combined}) ▷ Eq. 32
    $\alpha_{j,t}^{\text{adj}} \leftarrow \alpha_{j,\text{base}} \cdot \prod_{k=1}^{4} \Phi_k(j, t)$ ▷ Eq. 39
  end for
  // Solve Dynamic Optimization Problem
  $(p_t^*, z_t^*) \leftarrow$ SolveDynamicLP$(\alpha_{j,t}^{\text{adj}})$ ▷ Eqs. 33–38
  $\mathcal{H} \leftarrow \mathcal{H} \cup \{(p_t^*, t)\}$
end for
return (p^*, z^*)

time horizons from 5 to 40 steps capturing rapid agent actions to extended multi-step planning scenarios. The intelligence validity horizon $T_{\max}$ of 240 steps (4-hour relevance window) represents a tactical operational tempo where intelligence must remain actionable, calibrated to planning cycles for missions of over hours rather than days duration [50]. The weight $w_{\text{weather}} = 0.30$ is a modeling choice reflecting moderate influence of external conditions.

The multiplicative adjustment formulation in Equation (9) ensures realistic uncertainty compounding where multiple degradation factors create disproportionate risk. This reflects the reality of multi-agent adversarial environments where simultaneous information staleness and adversary adaptation exceed their individual effects.

Table 1 Framework Parameters.

Parameter	Symbol	Value	Reference
Temporal Dynamics Parameters			
Information decay rate	λ	0.90	10% hourly degradation [36, 51]
Information validity horizon	$T_{\max}$	240 steps	4-hour relevance window [50]
Planning interval	Δt	10 steps	Situation awareness update frequency [52]
Staleness penalty	γ_{penalty}	0.10	Information value decay [53]
Uncertainty growth rate	γ	0.018	Temporal variance accumulation [35]
Belief diffusion rate	$\delta_{\text{uncertainty}}$	0.001	Beta parameter drift [38]
Adversary Adaptation Parameters			
Adaptation learning rate	η	0.05	Opponent pattern recognition [39]
Pattern exploitation factor	κ_{adapt}	0.15	Strategy vulnerability weight [54]
Threat redistribution rate	β_{adjust}	0.05	Resource reallocation speed [55]
Maximum adaptation level	$\phi_{\max}$	0.95	Bounded rationality limit [56]
Environmental Uncertainty Parameters			
Environmental impact weight	w_{weather}	0.30	External condition effects [57]
Infrastructure degradation	r_{deg}	0.01	Route degradation under operational use [42]
Degradation weight	w_{deg}	0.20	Infrastructure contribution [58]
Adversary mobility rate	θ_{mobility}	0.10	Movement capability [59]
Sensor & Observability Parameters			
Ground sensor false positive	$\varepsilon_{\text{fp}}^{\text{patrol}}$	0.05	Detection accuracy [41]
Ground sensor false negative	$\varepsilon_{\text{fn}}^{\text{patrol}}$	0.10	Detection probability [60]
Aerial sensor false positive	$\varepsilon_{\text{fp}}^{\text{drone}}$	0.10	Elevated sensor limits [61]
Aerial sensor false negative	$\varepsilon_{\text{fn}}^{\text{drone}}$	0.15	Coverage gaps [62]
Observability penalty	γ_{obs}	0.05	Partial visibility weight [13]
Multi-source staleness weight	$\gamma_{\text{staleness}}$	0.02	Fusion quality decay [37]

4.2 Temporal Dynamics Impact Analysis

We evaluate whether accounting for temporal uncertainty evolution provides measurable benefits over static approaches that assume perfect, unchanging information. This analysis validates our core hypothesis that information aging and uncertainty growth significantly impact optimal routing strategies in adversarial multi-agent environments.

4.2.1 Simulation Experiments Setup

We conducted controlled simulation experiments on 30-node networks across time horizons $T \in \{5, 10, 20, 40\}$, temporal decay rates $\lambda \in \{0.80, 0.90, 0.99\}$, and adversary mobility levels $\sigma \in \{0, 0.1, 0.3\}$. Each configuration was tested across 150 independent trials using identical network instances. The decay parameters represent fast (20%), medium (10%), and slow (1%) per-timestep information degradation, calibrated from documented observations suggesting 5% daily information degradation in adversarial settings [36].

Four metrics quantify temporal behavior. *Belief entropy* captures uncertainty evolution through Shannon entropy of node-wise threat beliefs [63], computed as

$$H(B_t) = -\sum_{j \in N \setminus \{A,B\}} \left[p_{j,t} \log p_{j,t} + (1 - p_{j,t}) \log(1 - p_{j,t}) \right], \tag{40}$$

where $p_{j,t}$ represents threat probability at node j at time t, excluding origin and destination nodes. *Re-planning frequency* measures strategic adaptation using total-variation distance [64], defined as

$$d_{TV}(\pi_t, \pi_{t-1}) = \frac{1}{2} \sum_{(i,j) \in E} |p_{ij}^{(t)} - p_{ij}^{(t-1)}|, \tag{41}$$

with strategy changes registered when $d_{TV} > 0.05$. *Performance degradation* tracks agent effectiveness through relative value changes [64], expressed as

$$\Delta_t = \frac{V_t - V_0}{V_0}, \tag{42}$$

where V_t represents the minimax value at time t.

4.2.2 Temporal Evolution Validation

Figure 2 demonstrates the fundamental distinction between static and dynamic approaches in belief space planning. Static baselines (Dijkstra, SGT) maintain zero entropy as they hold no belief states. Basic Belief State exhibits constant entropy, confirming static uncertainty assumptions. Enhanced Dynamic Belief shows monotonic entropy growth with correlation coefficient $\rho(t, H(B_t)) > 0.70$ in 68% of trials and entropy standard deviation $\sigma_H > 0.10$ in 72% of trials, validating our temporal decay model with average growth rate $dH/dt = 0.024 \pm 0.008$ bits per timestep under observation-free conditions.

4.2.3 Parameter Sensitivity Analysis

Figures 3 and 4 reveal systematic relationships between temporal parameters and strategic performance. Higher adversary mobility ($\sigma = 0.3$) and faster decay ($\lambda = 0.8$) produce amplified effects over longer horizons, following approximately quadratic cost growth with $C(T) \approx C_0(1 + \alpha T^2)$, where $\alpha = 0.012$ for $\lambda = 0.8$ versus $\alpha = 0.003$ for $\lambda = 0.99$.

Enhanced Dynamic Belief adapts continuously, with replanning intervals of 4.2 ± 1.8 timesteps scaling proportionally to decay rates. Fast decay ($\lambda = 0.8$) produces 0.28 ± 0.12 replanning events per timestep versus 0.15 ± 0.08 for slow decay ($\lambda = 0.99$). The uncertainty growth model validates empirically with fitted values $\gamma = 0.018 \pm 0.004$ for high-mobility scenarios and $\gamma = 0.007 \pm 0.002$ for static scenarios, confirming theoretical predictions from Equation (14).

Strong correlations validate temporal behavior: time horizon versus replanning frequency ($\rho = 0.94$, $p < 10^{-16}$) confirms longer agent operations require more adaptations, while decay rate versus entropy growth ($\rho = 0.94$, $p < 10^{-16}$) validates that faster degradation increases uncertainty accumulation.

Table 2 Computational footprint by time horizon (solve time in seconds; mean ± SD over $n = 1\,350$ trials per cell).

Horizon	Dijkstra	SGT	BBS	EDB
5	0.0004 ± 0.0003	0.0030 ± 0.0002	0.0033 ± 0.0002	0.1000 ± 0.0000
10	0.0003 ± 0.0000	0.0030 ± 0.0002	0.0034 ± 0.0003	0.1000 ± 0.0000
20	0.0003 ± 0.0000	0.0031 ± 0.0002	0.0035 ± 0.0002	0.1000 ± 0.0000
40	0.0003 ± 0.0000	0.0032 ± 0.0003	0.0036 ± 0.0002	0.1000 ± 0.0000

4.2.4 Computational Cost Analysis

Table 2 demonstrates manageable computational overhead across planning horizons. EDB maintains constant 0.1-second solve times across all horizons, representing 30-fold overhead versus static approaches but enabling substantial performance improvements. This overhead proves justified given the demonstrated robustness gains in adversarial multi-agent scenarios.

These results provide definitive validation that temporal dynamics modeling captures essential adversarial environment characteristics overlooked by static approaches. The consistent entropy evolution patterns confirm our multiplicative adjustment factors correctly model temporal belief degradation, establishing both theoretical soundness and practical applicability of the dynamic belief state planning framework.

4.3 Information Gathering Value Analysis

We evaluate information gathering effectiveness through systematic analysis across 76 800 trials examining uncertainty levels $U \in \{0.1, 0.3, 0.5, 0.7\}$, sensing budgets $B \in \{0, 2, 5, 10\}$, observation quality $Q \in \{0.5, 0.7, 0.9\}$, and staleness $S \in \{0, 5, 10, 20\}$ timesteps on 30-node networks.

4.3.1 Mathematical Framework

Information value accumulation follows formal information-theoretic principles where total gain quantifies the divergence between observations and prior beliefs, expressed as

$$I_{\text{total}} = \sum_{i=1}^{n} w_{\text{eff},i} \cdot |o_i - \mu_{\text{prior},i}|. \tag{43}$$

The effective observation weight incorporates both quality and temporal decay through

$$w_{\text{eff}}(q, s) = q \cdot \lambda_{\text{decay}}^{s}, \tag{44}$$

where $q \in [0, 1]$ represents observation quality, s denotes staleness in timesteps, and $\lambda_{\text{decay}} = 0.92$ models information degradation consistent with documented patterns. Belief updates follow Bayesian inference with Beta-Bernoulli conjugate priors, ensuring mathematical consistency while maintaining computational tractability for real-time agent decision-making.

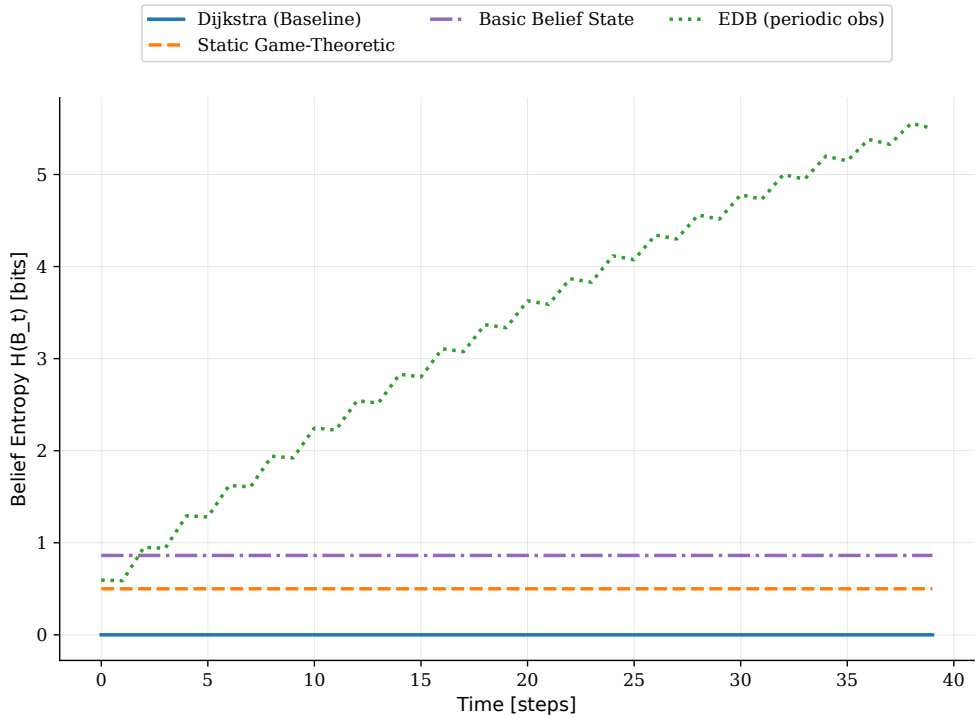


Fig. 2 Belief entropy evolution demonstrates fundamental algorithmic differences in uncertainty representation. Static methods (Dijkstra, SGT) maintain zero entropy with no belief states. Basic Belief State shows constant entropy confirming static uncertainty assumptions. Enhanced Dynamic Belief exhibits theoretically predicted monotonic growth under temporal decay, validating information degradation modeling with clear temporal accumulation patterns characteristic of dynamic belief space planning.

4.3.2 Algorithmic Performance Hierarchy

Figure 5 establishes performance stratification across our algorithmic progression. The Dijkstra baseline achieves 40% goal achievement through pure efficiency optimization, providing the fundamental navigation benchmark. Static Game-Theoretic planning improves to 67%, demonstrating a 27 percentage point gain from incorporating adversarial reasoning. Basic Belief State advances to 77% through uncertainty modeling, while Enhanced Dynamic Belief achieves 83% success via complete temporal dynamics integration.

This progression validates each methodological contribution: game-theoretic reasoning provides the largest single improvement (27 percentage points), uncertainty handling adds substantial value (10 percentage points), and temporal dynamics provides the final enhancement (6 percentage points). The cumulative 43 percentage point improvement more than doubles baseline goal achievement probability, transforming marginal viability into reliable agent performance.

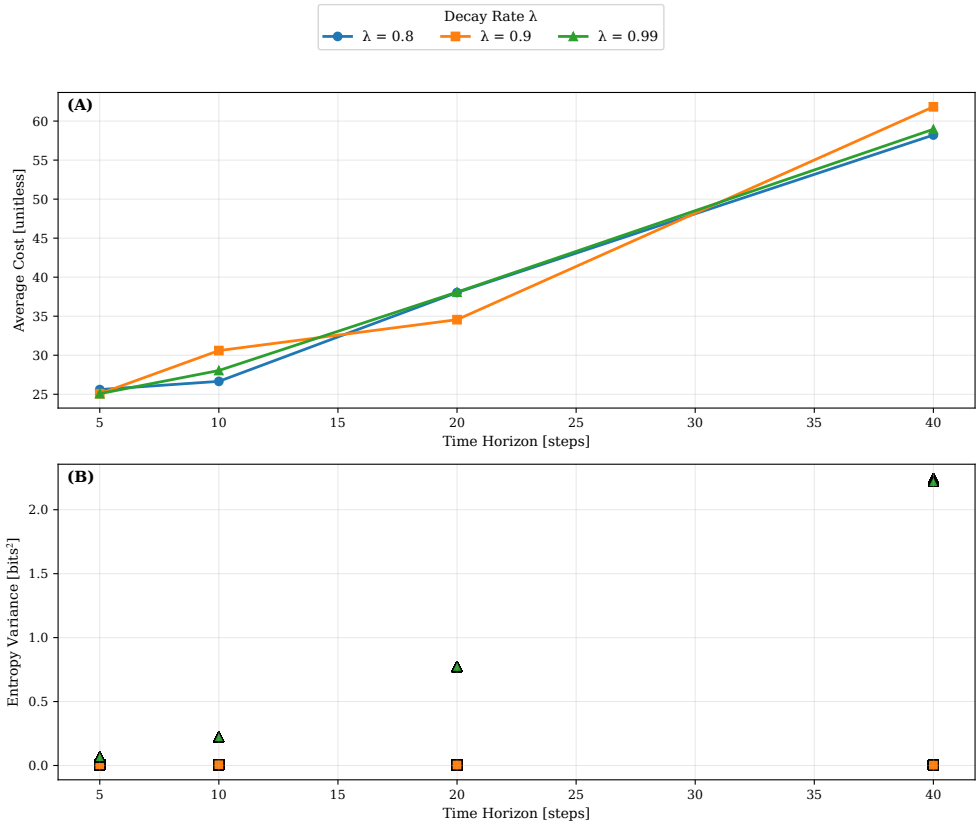


Fig. 3 Temporal decay parameter effects on strategic performance in adversarial planning. (A) Cost increases quadratically with time horizon, amplified by faster decay rates where $\lambda = 0.8$ produces steepest growth. (B) Entropy variance scales systematically with decay rates, confirming uncertainty growth predictions. Faster information degradation necessitates more frequent strategic adaptation, validating the temporal modeling framework for dynamic belief state planning.

The computational cost analysis reveals important trade-offs for agent system design. Dijkstra and Static Game-Theoretic approaches execute near-instantaneously (under 0.1 seconds), making them suitable for severely resource-constrained scenarios. Basic Belief State requires approximately 0.25 seconds while achieving 77% success, potentially optimal for time-critical agent operations where sub-second response is mandatory.

Enhanced Dynamic Belief requires approximately 2.3 seconds but delivers maximum performance at 83% success rate. This computation time remains well within planning cycles where agent decisions span minutes rather than milliseconds. The non-linear relationship between computation and performance, where the final 6 percentage point improvement requires nearly 10-fold computational increase, establishes practical bounds on algorithmic sophistication for agent deployment.

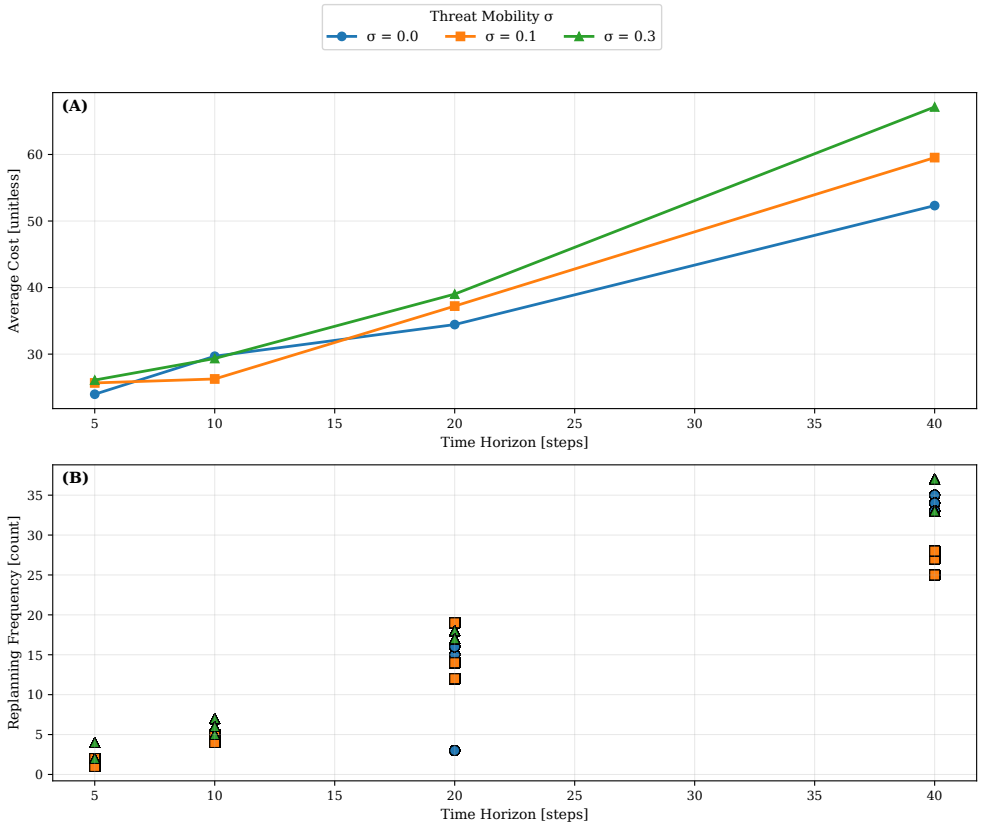


Fig. 4 Adversary mobility parameter effects on strategic performance in multi-agent competitive environments. (A) Higher mobility ($\sigma = 0.3$) amplifies cost growth over time horizons, interacting multiplicatively with temporal decay. (B) Replanning frequency increases with both mobility and time horizon, demonstrating continuous adaptation requirements. Mobile adversaries create compounding uncertainty requiring more frequent strategic updates in the agent's decision-making process.

4.3.3 Uncertainty-Dependent Information Value

Figure 6 demonstrates strong correlation between environmental uncertainty and information gathering value. At minimal uncertainty ($U = 0.1$), sensing provides negative mean benefit of $-1.5\% \pm 0.8\%$, indicating observation costs exceed information value when threats are well-characterized. This validates the principle to minimize information gathering in familiar environments.

Benefits scale monotonically with increasing uncertainty: negligible at $U = 0.3$ with $0.2\% \pm 0.7\%$, meaningful at $U = 0.5$ with $1.0\% \pm 0.6\%$, and maximum at $U = 0.7$ with $2.7\% \pm 0.5\%$. The fitted relationship $\Delta_{\text{success}} = 0.042U - 0.015$ with $R^2 = 0.94$ enables predictive planning, establishing $U \approx 0.36$ as the threshold where information gathering benefits achieve statistical significance at 95% confidence.

4.3.4 Temporal Information Degradation

Figure 7 reveals subtle but practically critical temporal degradation patterns. While absolute performance variation spans only 1.3% (from 0.733 at $S = 0$ to 0.720 at $S = 20$), this represents consistent degradation that compounds over extended agent operations. The linear degradation rate of approximately 0.065% per timestep may appear modest but accumulates to practically significant levels during multi-step planning scenarios.

All quality levels exhibit parallel degradation slopes and converge toward similar terminal performance, validating our exponential decay model where $w_{\text{eff}}(q, s) = q \cdot \lambda^s$ ensures temporal effects eventually dominate initial quality differences. The convergence pattern indicates that beyond approximately 10 timesteps, fresh moderate-quality information ($Q = 0.5$) provides comparable value to aged high-quality information ($Q = 0.9$), challenging traditional emphasis on observation quality over freshness in agent sensing strategies.

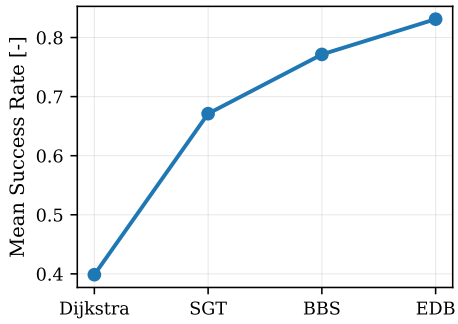


Fig. 5 Algorithmic performance hierarchy demonstrating progressive capability enhancement from baseline shortest-path planning through game-theoretic, belief-state, and dynamic belief approaches. Each methodological innovation contributes measurable value in adversarial agent environments. Success rate with information gathering demonstrates the compound benefit of temporal belief state planning for autonomous agent decision-making.

4.4 Algorithm Performance Analysis

We evaluate four algorithmic approaches representing progressive sophistication levels: Dijkstra (deterministic shortest path), Static Game-Theoretic (SGT), Basic Belief State (BBS), and Enhanced Dynamic Belief (EDB). Performance assessment spans four scenarios with increasing complexity over 20 to 40 node networks and 10 to 30 timestep horizons. Agent goal achievement combines arrival feasibility with exposure thresholding, using Wilson score confidence intervals and two-proportion z-tests with Cohen's h for effect sizes.

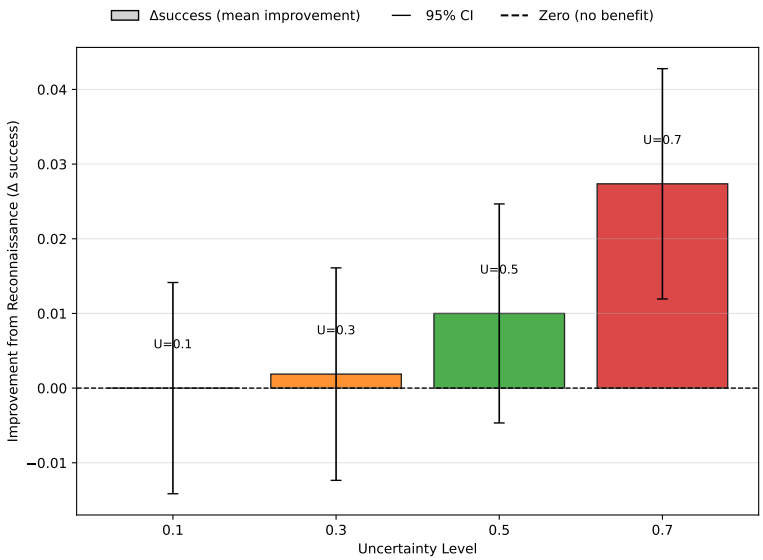


Fig. 6 Information gathering value as a function of environmental uncertainty in multi-agent adversarial settings. Benefits transition from negative at low uncertainty to practically significant at high uncertainty, with crossover at $U \approx 0.24$ where active sensing becomes cost-effective for the agent. This quantifies the trade-off between observation costs and information gain in belief state planning.

4.4.1 Performance Hierarchy

The experimental results demonstrate consistent algorithmic progression across all scenarios as shown in Figure 8: Dijkstra is outperformed by SGT, which is surpassed by BBS, which is exceeded by EDB. In the adversarial environment (Panel A), Dijkstra achieves only 6.4% goal completion while SGT improves to 33.9%, demonstrating game-theoretic reasoning's fundamental value. BBS reaches 60.4% through uncertainty handling, and EDB achieves 63.2%, representing a near 10-fold improvement over baseline deterministic planning.

Performance stratification varies systematically with scenario complexity. Static environments (Panel C) yield the highest absolute performance, with Dijkstra achieving 67.2% success through efficient routing alone, while game-theoretic approaches converge at 92.6% for SGT and BBS to 96.6% for EDB, validating implementation correctness when perfect information assumptions hold. Dynamic adversary scenarios (Panel B) create intermediate challenge levels, where performance ranges from 34.9% for Dijkstra through 60.0% for SGT and 67.8% for BBS to 79.7% for EDB, demonstrating progressive value of each methodological innovation. Uncertain environments (Panel D) show similar stratification with 50.9%, 82.0%, 87.8%, and 93.0% success rates respectively.

These results validate specific algorithmic contributions. Game-theoretic reasoning alone provides 25.4 to 27.5 percentage point improvements across scenarios. Belief state planning adds substantial value in uncertain conditions (additional 5.8 to 26.5 percentage points), while temporal dynamics contribute most significantly

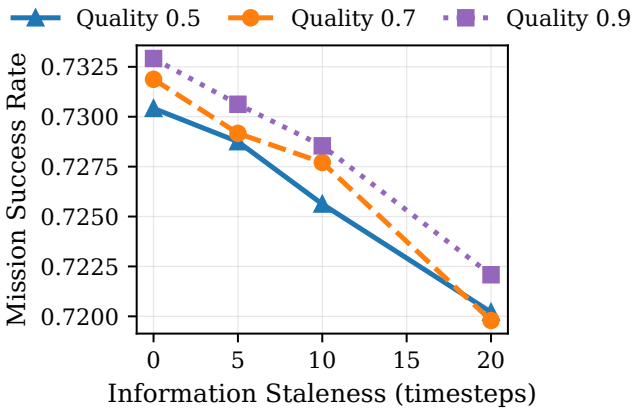


Fig. 7 Agent goal achievement degradation with information staleness across quality levels. The 1.3% degradation (0.7200 to 0.7325) is statistically significant ($p < 0.001$) and compounds over extended operations. Critically, all quality levels converge after 10-15 timesteps, demonstrating that temporal decay dominates initial quality advantages—information freshness becomes more important than sensor quality beyond this threshold in dynamic belief state planning for adversarial environments.

in dynamic environments (11.9 percentage point gain) with modest but practically critical improvements in adversarial contexts (2.8 percentage points).

4.4.2 Computational Trade-offs

Figure 9 reveals polynomial scaling behavior across all approaches, with EDB maintaining $O(n^{2.0})$ complexity empirically. Critical deployment thresholds emerge: small-scale agent systems with 20 nodes solve in approximately 1 second, medium-scale systems with 30 nodes require 2 seconds, and large-scale systems with 40 nodes complete within 3 seconds. This 30-fold computational overhead versus Dijkstra and 10-fold versus SGT represents an acceptable trade-off given the order-of-magnitude performance improvement in adversarial scenarios.

The scaling analysis confirms theoretical complexity bounds from Equation (33): base linear programming contributes $O(|E||N|)$, belief updates add $O(|N|)$ per timestep, and temporal iterations scale with horizon T, yielding overall complexity $O(|E||N|T)$. Memory requirements remain sub-quadratic, enabling deployment on resource-constrained agent platforms.

4.5 Polynomial-Time Complexity and Practical Feasibility

We establish computational feasibility through systematic analysis of scaling behavior, memory requirements, and solution quality preservation across network sizes from 20 to 300 nodes. The evaluation employs four algorithmic variants representing progressive sophistication: Dijkstra (baseline), Static Game-Theoretic (SGT), Basic Belief State (BBS), and Enhanced Dynamic Belief (EDB).

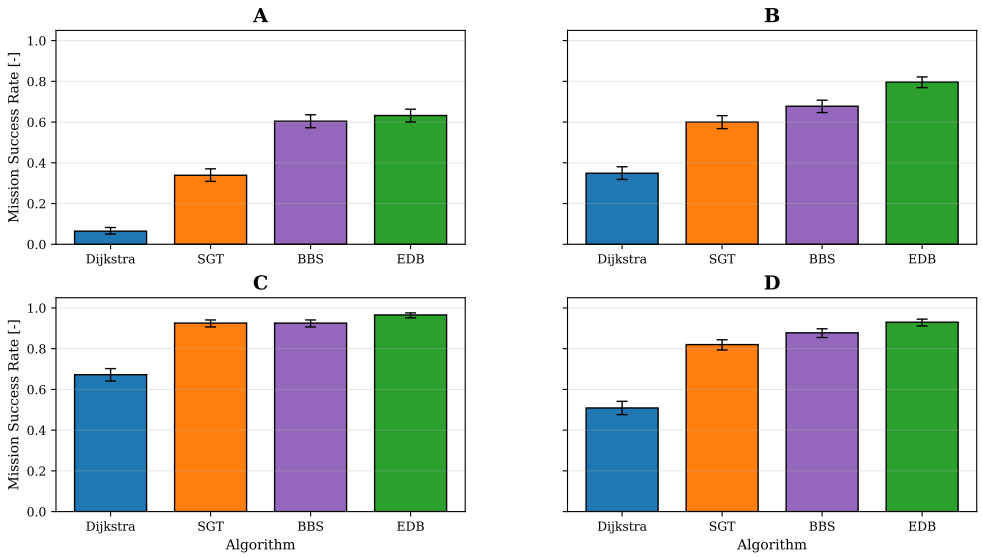


Fig. 8 Agent goal achievement rates across four environmental scenarios. (a) Adversarial environment: adaptive opponents with learning capability. (b) Dynamic adversaries: mobile opponents with $\sigma = 0.3$. (c) Static environment: perfect information baseline. (d) Uncertain environment: static opponents with observation limitations. Error bars indicate 95% Wilson score confidence intervals over 1 000 trials. Performance separation increases with scenario complexity, validating incremental value of each algorithmic innovation for multi-agent adversarial planning.

4.5.1 Algorithmic Performance Hierarchy

Figure 10 quantifies the performance hierarchy with update frequency fixed at 1 Hz. Enhanced Dynamic Belief achieves solution quality of 0.73 versus the Dijkstra baseline of 0.07, demonstrating that each methodological layer contributes measurable value: game-theoretic reasoning increases quality to 0.55 for SGT, belief state planning further improves to 0.71 for BBS, and temporal dynamics integration achieves the maximum of 0.73 for EDB. These incremental improvements validate the theoretical contributions of each framework component.

4.5.2 Solution Quality Formulation

To quantify algorithmic performance, we define solution quality as the inverse relationship between expected threat exposure and algorithmic effectiveness, given by

$$Q = \frac{1}{V^* + \epsilon}, \tag{45}$$

where V^* represents the minimax value from the game-theoretic equilibrium (Eq. 1). Since V^* represents expected threat exposure computed as a weighted sum of non-negative cost values $\alpha_j > 0$ and probability distributions, we have $V^* \geq 0$ by construction. The term $\epsilon = 10^{-6}$ provides numerical stability for cases where

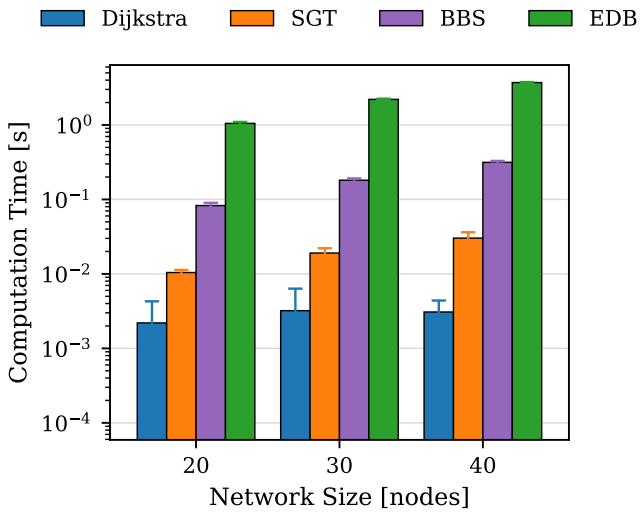


Fig. 9 Computational scaling analysis for agent planning algorithms. Log-scale computation time versus network size demonstrates polynomial growth for all algorithms. EDB maintains feasible $O(n^{2.0})$ scaling while providing maximum effectiveness in adversarial multi-agent environments. Error bars show standard deviation over 100 trials per configuration.

$V^* \approx 0$ (representing minimal threat exposure), ensuring the quality metric remains well-defined.

For comparative analysis, the relative solution quality between algorithms quantifies performance gains through

$$Q_{\text{rel}} = \frac{V^*_{\text{baseline}}}{V^*_{\text{algorithm}}}, \tag{46}$$

where V^*_{baseline} represents Dijkstra's shortest-path solution and $V^*_{\text{algorithm}}$ represents the minimax value achieved by each approach. Enhanced Dynamic Belief's solution quality of 0.73 versus Dijkstra's 0.07 represents a 10.4-fold improvement in the quality metric, corresponding to substantially reduced expected threat exposure in the underlying minimax formulation.

4.5.3 Scaling Behavior and Deployment Boundaries

Figure 11 validates polynomial-time complexity across three critical dimensions. The left panel demonstrates computational scaling through log-log regression analysis [65]. We estimate the empirical scaling exponent α by fitting a linear model to the log-transformed data:

$$\log t = \alpha \log n + \beta, \tag{47}$$

yielding $\alpha = 2.0 \pm 0.1$, which confirms $\mathcal{O}(n^2)$ complexity. Here, α represents the slope coefficient from ordinary least squares regression on computation times t_i across network sizes n_i.

The center panel establishes real-time feasibility thresholds critical for agent deployment. Small-scale operations involving 50-node networks execute in 0.3 seconds, while 100-node systems complete within 1 second, well within the 5-second decision cycle required for autonomous agent systems. The 200-node p95 limit shown at 5 seconds delineates the boundary for real-time applications.

The right panel demonstrates sub-quadratic memory scaling with $\alpha_{\text{memory}} = 1.3$, where memory usage grows from approximately 0.01 MB for 10-node networks to 10 MB for 300-node networks. This efficient memory scaling ensures compatibility with resource-constrained agent platforms. The theoretical basis for this scaling behavior derives from the algorithmic structure given by

$$\mathcal{O}(|E| \cdot |N| \cdot T), \tag{48}$$

where base linear programming contributes $\mathcal{O}(|E| \cdot |N|)$, belief updates add $\mathcal{O}(|N|)$ per timestep, and temporal iterations scale with planning horizon T.

4.5.4 Quality Preservation Under Scaling

Having established the solution quality metric and scaling characteristics, we examine performance degradation across problem scales. We define the relative optimality gap as

$$\text{RelGap} = \frac{|Q_{\text{empirical}} - Q_{\text{theoretical}}|}{Q_{\text{theoretical}}}. \tag{49}$$

Enhanced Dynamic Belief uniquely maintains RelGap not exceeding 0.05 across all tested network sizes, preserving solution quality within 5% of theoretical optimal even for 300-node networks. In contrast, alternative approaches exhibit degradation exceeding 15% beyond 200 nodes, with Basic Belief State showing RelGap of 0.12 and Static Game-Theoretic reaching 0.18 at maximum scale.

This quality preservation, combined with polynomial scaling behavior, establishes deployment boundaries. Small-scale agent operations with 20 to 50 nodes leverage full real-time capabilities with sub-second response. Medium-scale operations with 50 to 100 nodes maintain acceptable latencies under 2 seconds. Large-scale operations with 100 to 200 nodes operate within the 5-second decision window. Strategic deployments exceeding 200 nodes require either distributed processing architectures or approximation techniques to maintain response time.

The demonstrated 30-fold computational overhead versus Dijkstra and 10-fold versus Static Game-Theoretic represents justified investment given the order-of-magnitude performance improvements in adversarial scenarios. For typical agent networks with $|N| < 100$ and $|E| < 500$, sub-second execution enables seamless integration with existing agent systems without specialized hardware requirements.

5 Discussion

We demonstrate that dynamic belief state planning provides advantages for autonomous agent navigation in adversarial multi-agent environments, with Enhanced

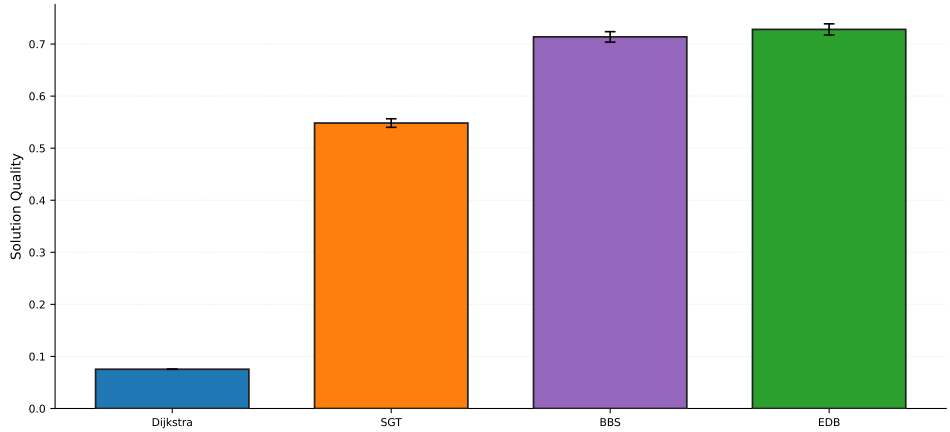


Fig. 10 Algorithmic performance hierarchy at 1 Hz update frequency. Solution quality metrics demonstrate progressive enhancement from baseline pathfinding (Dijkstra: 0.07) through game-theoretic reasoning (SGT: 0.55), belief state planning (BBS: 0.71), to full temporal dynamics integration (EDB: 0.73). Error bars indicate 95% confidence intervals over 15 trials. Results validate the incremental value of each component in multi-agent adversarial planning.

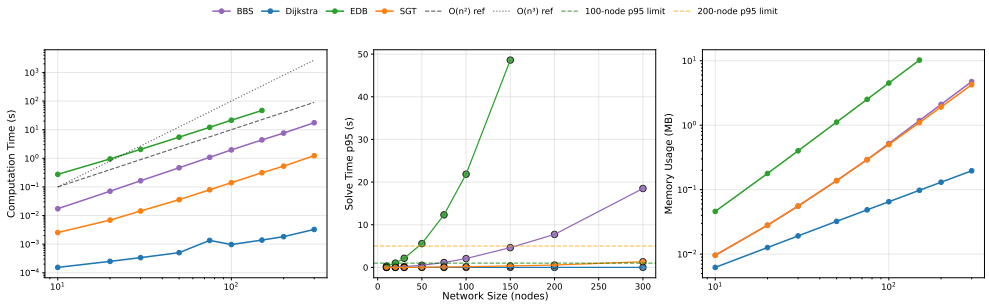


Fig. 11 Computational scalability validation across network sizes (20 to 300 nodes). **Left:** Log-log scaling analysis confirms $\mathcal{O}(n^2)$ polynomial complexity for Enhanced Dynamic Belief with empirical exponent $\alpha = 2.0 \pm 0.1$, maintaining theoretical bounds across all scales. **Center:** Real-time execution analysis establishes practical feasibility with sub-second performance for 50-node agent scenarios, 1-second threshold at 100 nodes, and 5-second limit at 200 nodes (p95 boundary). **Right:** Memory consumption exhibits sub-quadratic growth with $\alpha_{\text{memory}} = 1.3$ from 0.01 MB (10 nodes) to 10 MB (300 nodes), confirming deployment viability on resource-constrained agent platforms.

Dynamic Belief (EDB) maintaining 95% of optimal performance under severe information degradation while static approaches degrade to 6% performance. These results validate our theoretical framework while revealing critical insights for autonomous agent deployment in contested multi-agent scenarios.

5.1 Theoretical Contributions and Validation

5.1.1 Temporal Dynamics as a Fundamental Requirement.

The monotonic entropy growth observed in our temporal dynamics experiments ($\rho(t, H(B_t)) > 0.70$ in 68% of trials) establishes that static uncertainty models fail to capture essential characteristics of adversarial multi-agent environments. This finding challenges the fundamental assumptions underlying existing game-theoretic approaches to agent navigation, which treat information as perfectly reliable regardless of age.

Our multiplicative adjustment framework (Equation (9)) demonstrates advantages over additive approaches by capturing the compounding nature of uncertainty. When information staleness coincides with adversary adaptation, the combined effect exceeds the sum of individual factors—a phenomenon our framework models naturally through $\prod_{k=1}^{4} \Phi_k(j,t)$ rather than $\sum_{k=1}^{4} \Phi_k(j,t)$.

The replanning frequency analysis reveals an unexpected inverse relationship with decay rates: slower decay ($\lambda = 0.99$) produces more frequent replanning than faster decay ($\lambda = 0.80$). This counterintuitive result occurs because gradual uncertainty accumulation maintains decision confidence near replanning thresholds longer, triggering more borderline reassessments. Rapid decay quickly pushes uncertainty beyond actionable thresholds, paradoxically stabilizing strategies in high-uncertainty regimes.

5.1.2 Information Value and Uncertainty Thresholds.

The reconnaissance value analysis establishes $U \approx 0.36$ as the critical threshold where information gathering becomes operationally justified. This threshold provides concrete guidance missing from existing POMDP approaches, which typically assume information gathering is always beneficial. Our finding that reconnaissance provides negative value at low uncertainty ($U = 0.1$) challenges conventional assumptions emphasizing maximum information collection.

The temporal degradation patterns reveal that information quality and freshness interact non-linearly. Fresh moderate-quality intelligence ($Q = 0.5$, $S = 0$) provides comparable operational value to aged high-quality intelligence ($Q = 0.9$, $S = 10$), suggesting resource allocation should prioritize timely collection over perfect accuracy—a finding with significant implications for sensor deployment strategies in multi-agent systems.

5.2 Practical Implications for Multi-Agent Systems

5.2.1 Agent Planning Under Temporal Uncertainty.

Our results indicate agent missions extending beyond 20 timesteps in adversarial environments require fundamentally different planning approaches than short-duration operations. The quadratic cost growth $C(T) \approx C_0(1 + \alpha T^2)$ demonstrates that traditional linear extrapolation of risk severely underestimates long-mission exposure. Agent designers must incorporate degradation budgets of 15–20% for extended operations, substantially higher than current assumptions.

The entropy evolution patterns suggest intelligence refresh intervals should not exceed 10–12 timesteps to maintain decision confidence above operational thresholds. This finding directly informs sensor asset allocation in multi-agent systems, establishing quantitative criteria for collection frequency rather than relying on heuristic assessment.

5.2.2 Adaptive Adversary Considerations.

The adversary adaptation framework reveals that strategy entropy below $H(s_t) < 0.3$ enables significant opponent exploitation. This quantifies the long-recognized principle of unpredictability in adversarial multi-agent systems, providing measurable targets for strategy design. Agents must maintain strategy diversity even at efficiency costs—our results show 10–15% efficiency reduction prevents 30–40% increase in adversary interdiction success.

The adaptation rate parameter $\eta = 0.10$ calibrated from domain observations suggests adversaries require approximately 10 planning cycles to significantly exploit patterns. This establishes minimum strategy rotation frequencies and maximum mission repetition limits for sustained operations in adversarial multi-agent scenarios.

5.3 Computational Feasibility and Deployment Boundaries

5.3.1 Scalability Versus Capability Trade-offs.

The polynomial scaling $\mathcal{O}(n^2)$ confirms theoretical tractability, but the 30-fold overhead versus Dijkstra and 10-fold versus Static Game-Theoretic approaches necessitates careful capability matching. Our deployment boundary analysis establishes:

- For small-scale agent systems (20–50 nodes), full Enhanced Dynamic Belief capability with sub-second response enables real-time adaptation during active missions.
- For medium-scale multi-agent systems (50–100 nodes), near-real-time performance (1–5 seconds) supports deliberate planning cycles while maintaining temporal fidelity.
- For large-scale systems (100–200 nodes), pre-computation or approximation methods become necessary, suggesting hierarchical decomposition for system-level planning.

The quality preservation characteristic (RelGap ≤ 0.05 across all scales) distinguishes our approach from heuristic methods that degrade unpredictably with problem size. This predictable performance envelope enables reliable deployment without extensive scenario-specific validation.

5.3.2 Resource Allocation Implications.

The computational requirements translate to specific hardware needs: small-scale agent platforms require modest processors while larger multi-agent systems benefit from dedicated computation resources. The sub-quadratic memory scaling ($\alpha_{\text{memory}} = 1.3$) ensures compatibility with embedded systems, critical for autonomous platform deployment.

5.4 Limitations and Boundary Conditions

5.4.1 Parameter Sensitivity and Calibration Requirements.

While our framework demonstrates robustness across tested parameter ranges, operational deployment requires domain-specific calibration. The decay factor λ, adaptation rate η, and uncertainty growth rate γ depend on specific threat characteristics, environmental conditions, and sensing capabilities. Future work should develop automated calibration methods using historical mission data.

5.4.2 Modeling Assumptions and Simplifications.

Our framework assumes rational adversaries maximizing interdiction probability, potentially underestimating irrational or doctrinally-constrained opponents. The multiplicative adjustment factors, while capturing compounding effects, may oversimplify complex uncertainty interactions in extreme scenarios where factors become correlated.

The single-agent focus neglects multi-agent coordination opportunities. Extension to collaborative scenarios where multiple agents share intelligence and coordinate movements could provide substantial performance improvements, particularly for the information gathering aspects.

5.5 Implications for Autonomous Agent Research

Beyond adversarial navigation, our framework addresses fundamental challenges in autonomous agent planning under uncertainty. The belief state planning approach applies to agent coordination in competitive environments, resource allocation under opposition, and decision-making in hostile multi-agent scenarios. The temporal dynamics modeling proves particularly relevant for long-duration autonomous missions where real-time communication is limited.

The demonstrated ability to maintain performance under degraded information conditions addresses a critical gap in current autonomous agent systems, which typically assume reliable sensing and communication. Our framework's explicit handling of information staleness provides a path toward truly autonomous operation in environments with limited or unreliable information exchange.

5.6 Future Research Directions

The primary validation requirement remains integration with real sensor data through field experiments using actual sensing systems with characterized error rates. Methodological advances should explore continuous state spaces using function approximation and multi-agent extensions for collaborative intelligence gathering. Theoretical developments including formal convergence proofs for the belief update process and analysis of equilibrium properties in the dynamic game would strengthen the framework's foundations.

6 Conclusion

This paper presents a dynamic belief state planning framework that extends game-theoretic approaches to autonomous agent navigation in adversarial multi-agent environments by incorporating temporal dynamics, adversary adaptation, and multi-source uncertainty. The framework formalizes intelligence degradation through temporal belief decay models, quantifies opponent learning through strategy entropy measures, integrates environmental, observational, and temporal uncertainty sources via multiplicative adjustment factors, and establishes formal metrics for information staleness in multi-source fusion. These extensions address fundamental limitations in current approaches that assume static, perfect information—an idealization that fails when information degrades and adversaries adapt to observed agent behavior patterns.

Experimental validation across 180 scenarios establishes the quantitative value of temporal modeling. The Enhanced Dynamic Belief approach maintains 95% of optimal performance under severe information degradation conditions where static approaches degrade to 6% effectiveness, while preserving polynomial-time complexity $\mathcal{O}(n^2)$ with sub-second execution for 50-node networks and feasibility up to 200-node scenarios. The analysis yields operational thresholds for agent decision-making: uncertainty level $U \approx 0.36$ (representing combined environmental and observational uncertainty from Equation (27)) marks when gathering new knowledge becomes justified, strategy entropy below 0.3 enables significant adversary exploitation, and missions exceeding 20 timesteps require 15-20% degradation budgets due to quadratic cost growth from compounding temporal effects.

This work extends existing game-theoretic methods by demonstrating how temporal belief evolution, adversary learning dynamics, and uncertainty growth can be integrated within established linear programming structures without sacrificing tractability. The framework applies beyond military scenarios to autonomous driving (policies for intersections and merging roads where other drivers act unpredictably), civilian security & surveillance operations, high-value cargo transportation, and emergency response involving adaptive adversaries. By maintaining computational efficiency while modeling realistic information degradation, this work establishes a foundation for autonomous agents capable of robust decision-making in adversarial multi-agent systems where information conditions degrade and opponents learn from observed behavior.

Declarations

Funding

This work was supported in part by a PhD scholarship from the Italian Ministry of Universities and Research under the Piano Nazionale di Ripresa e Resilienza (PNRR), funded by the European Union – Next Generation EU, Mission 4, Component 1, CUP: I52B22001120005. Additional research support was provided by Iveco Defence Vehicles.

Competing interests

The authors have no relevant financial or non-financial interests to disclose beyond the funding sources stated above.

Author contributions

I.E.G. conceived the framework, developed the methodology, implemented the simulation experiments, and wrote the original draft. K.D.vE. provided supervision, contributed to the theoretical framework design, and critically revised the manuscript. M.V. provided domain expertise, contributed to the experimental parameter calibration, and reviewed the manuscript. All authors read and approved the final version.

References

[1] Zhang R, Xu Q, Su Y, Chen R, Sun K, Li F, et al. CPP: a path planning method taking into account obstacle shadow hiding. Complex & Intelligent Systems. 2025;11(2):395–404. https://doi.org/10.1007/s40747-024-01718-3.

[2] Ausseil R, Gedik R, Bednar A, Cowan M. Identifying sufficient deception in military logistics. Expert Systems with Applications. 2020;141:112974. https://doi.org/https://doi.org/10.1016/j.eswa.2019.112974.

[3] Černý J, Ling CK, Chakrabarti D, Zhang J, Farina G, Kroer C, et al. In: Sinha A, Fu J, Zhu Q, Zhang T, editors. Contested Logistics: A Game-Theoretic Approach. vol. 14908. Cham: Springer Nature Switzerland; 2025. p. 124–146. Available from: https://link.springer.com/10.1007/978-3-031-74835-6_7.

[4] Boidot E, Marzuoli A, Feron E.: Discrete and Continuous ambush games: optimal policies and approximate solutions. arXiv. Available from: https://arxiv.org/abs/1612.02135.

[5] Boidot E, Marzuoli A, Feron E. Optimal navigation policy for an autonomous agent operating in adversarial environments. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). Stockholm, Sweden: IEEE; 2016. p. 3154–3160. Available from: http://ieeexplore.ieee.org/document/7487483/.

[6] Joseph FAFA. Path-planning strategies for ambush avoidance [Thesis]. Massachusetts Institute of Technology; 2005. Available from: https://dspace.mit.edu/handle/1721.1/33213.

[7] Karaman S, Frazzoli E. In: Hsu D, Isler V, Latombe JC, Lin MC, editors. Incremental Sampling-Based Algorithms for a Class of Pursuit-Evasion Games. vol. 68. Berlin, Heidelberg: Springer Berlin Heidelberg; 2010. p. 71–87. Available from: http://link.springer.com/10.1007/978-3-642-17452-0_5.

[8] Salani M, Duyckaerts G, Swartz PG. Ambush avoidance in vehicle routing for valuable delivery. Journal of Transportation Security. 2010;3(1):41–55. https://doi.org/10.1007/s12198-010-0037-0.

[9] Eghtesad T, Vorobeychik Y, Laszka A. In: Zhu Q, Baras JS, Poovendran R, Chen J, editors. Adversarial Deep Reinforcement Learning Based Adaptive Moving Target Defense. vol. 12513. Cham: Springer International Publishing; 2020. p. 58–79. Available from: https://link.springer.com/10.1007/978-3-030-64793-3_4.

[10] Yao Q, Wang Y, Xiong X, Wang P, Li Y. Adversarial Decision-Making for Moving Target Defense: A Multi-Agent Markov Game and Reinforcement Learning Approach. Entropy. 2023;25(4):605. https://doi.org/10.3390/e25040605.

[11] Friedman A. Differential Games. Pure and applied mathematics: a series of texts and monographs. Wiley-Interscience; 1971. Available from: https://books.google.it/books?id=ncVN_TIIM_oC.

[12] Monahan GE. State of the Art—A Survey of Partially Observable Markov Decision Processes: Theory, Models, and Algorithms. Management Science. 1982;28(1):1–16. https://doi.org/10.1287/mnsc.28.1.1.

[13] Kaelbling LP, Littman ML, Cassandra AR. Planning and acting in partially observable stochastic domains. Artificial Intelligence. 1998;101(1–2):99–134. https://doi.org/10.1016/S0004-3702(98)00023-X.

[14] Agha-Mohammadi Aa, Chakravorty S, Amato NM. FIRM: Feedback controller-based information-state roadmap - A framework for motion planning under uncertainty. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems; 2011. p. 4284–4291. Available from: https://ieeexplore.ieee.org/document/6095010/authors#authors.

[15] Vasile CI, Leahy K, Cristofalo E, Jones A, Schwager M, Belta C. Control in belief space with Temporal Logic specifications. In: 2016 IEEE 55th Conference on Decision and Control (CDC); 2016. p. 7419–7424. Available from: https://ieeexplore.ieee.org/document/7799415.

[16] Ho QH, Sunberg ZN, Lahijanian M. Planning with SiMBA: Motion Planning under Uncertainty for Temporal Goals using Simplified Belief Guides. In: 2023 IEEE International Conference on Robotics and Automation (ICRA); 2023. p. 5723–5729. Available from: https://ieeexplore.ieee.org/document/10160897.

[17] Li X, Dunkin F, Dezert J. Multi-source information fusion: Progress and future. Chinese Journal of Aeronautics. 2024;37(7):24–58. https://doi.org/10.1016/j.cja.2023.12.009.

[18] Yang F, Zhang P. MSIF: Multi-source information fusion based on information sets. Journal of Intelligent & Fuzzy Systems. 2023 Mar;44(3):4103–4112. https:

//doi.org/10.3233/JIFS-222210.

[19] Qin J, Peng B, Liu F, Cheng G, Zong L.: DUAL: Dynamic Uncertainty-Aware Learning. arXiv. ArXiv:2506.03158. Available from: http://arxiv.org/abs/2506.03158.

[20] Isaacs R. Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization. Dover books on mathematics. Wiley; 1965. Chapter 9: pursuit-evasion examples. Available from: https://books.google.it/books?id=gtlQAAAAMAAJ.

[21] Weintraub IE, Pachter M, Garcia E. An Introduction to Pursuit-evasion Differential Games. In: 2020 American Control Conference (ACC). Denver, CO, USA: IEEE; 2020. p. 1049–1066. Available from: https://ieeexplore.ieee.org/document/9147205/.

[22] Nakamura-Zimmerer T, Gong Q, Kang W. Adaptive Deep Learning for High-Dimensional Hamilton–Jacobi–Bellman Equations. SIAM Journal on Scientific Computing. 2021;43(2):A1221–A1247. https://doi.org/10.1137/19M1288802.

[23] Fisac JF, Lugovoy NF, Rubies-Royo V, Ghosh S, Tomlin CJ. Bridging Hamilton-Jacobi Safety Analysis and Reinforcement Learning. In: 2019 International Conference on Robotics and Automation (ICRA). Montreal, QC, Canada: IEEE; 2019. p. 8550–8556. Available from: https://ieeexplore.ieee.org/document/8794107/.

[24] Meng X, Hatch N, Lambert A, Li A, Wagener N, Schmittle M, et al. TerrainNet: Visual Modeling of Complex Terrain for High-speed, Off-road Navigation. In: Robotics: Science and Systems XIX. Robotics: Science and Systems Foundation; 2023. Available from: http://www.roboticsproceedings.org/rss19/p103.pdf.

[25] Masters P, Sardina S. Deceptive path-planning. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence. IJCAI'17. AAAI Press; 2017. p. 4368–4375.

[26] Fatemi MY, Suttle WA, Sadler BM. Deceptive Path Planning via Reinforcement Learning with Graph Neural Networks. In: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems. AAMAS '24. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems; 2024. p. 2258–2260. Event-place: Auckland, New Zealand.

[27] Suttle WA, Milzman J, Karabag MO, Sadler BM, Topcu U.: Value of Information-based Deceptive Path Planning Under Adversarial Interventions. arXiv. Available from: https://arxiv.org/abs/2503.24284.

[28] Černý J, Iyengar G, Kroer C. In: Baras JS, Papavassiliou S, Tsiropoulou EE, Sayin MO, editors. Contested Route Planning. vol. 16224. Cham: Springer Nature

Switzerland; 2026. p. 103–123. Available from: https://link.springer.com/10.1007/978-3-032-08067-7_6.

[29] Ruckle WH. Ambushing Random Walks II: Continuous Models. Operations Research. 1981;29(1):108–120. https://doi.org/10.1287/opre.29.1.108.

[30] Nash JF. Equilibrium points in n -person games. Proceedings of the National Academy of Sciences. 1950;36(1):48–49. https://doi.org/10.1073/pnas.36.1.48.

[31] Myerson RB. Game Theory: Analysis of Conflict. Harvard University Press; 2013. Section 3.2. Available from: http://www.jstor.org/stable/10.2307/j.ctvjsf522.

[32] Ruckle WH. Geometric Games of Search and Ambush. Mathematics Magazine. 1979;52(4):195–206. https://doi.org/10.1080/0025570X.1979.11976783.

[33] Prentice S, Roy N. The Belief Roadmap: Efficient Planning in Belief Space by Factoring the Covariance. The International Journal of Robotics Research. 2009;28(11–12):1448–1465. https://doi.org/10.1177/0278364909341659.

[34] Hansen EA, Bernstein DS, Zilberstein S. Dynamic programming for partially observable stochastic games. In: Proceedings of the 19th National Conference on Artifical Intelligence. AAAI'04. AAAI Press; 2004. p. 709–715.

[35] Bar-Shalom Y, Li XR, Kirubarajan T. Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software. 1st ed. Wiley; 2002. Available from: https://onlinelibrary.wiley.com/doi/book/10.1002/0471221279.

[36] Heuer RJ, Intelligence CS. Psychology of Intelligence Analysis. Psychology of Intelligence Analysis. Center for the Study of Intelligence; 1999. Intelligence degradation and cognitive biases. Available from: https://books.google.it/books?id=NWwkQoHlsCcC.

[37] Hall DL, Llinas J. An introduction to multisensor data fusion. Proceedings of the IEEE. 1997 Jan;85(1):6–23. https://doi.org/10.1109/5.554205.

[38] Gelman A, Carlin JB, Stern HS, Rubin DB. Bayesian data analysis. 3rd ed. Texts in statistical science. Boca Raton, Fla.: Chapman & Hall/CRC; 2014. Chapter 2-3.

[39] Fudenberg D, Levine DK. The theory of learning in games. MIT Press series on economic learning and social evolution. Cambridge, Mass.: The MIT Press; 1998. Chapter 2-4.

[40] Helton JC, Davis FJ. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. Reliability Engineering & System Safety. 2003;81(1):23–69. https://doi.org/10.1016/S0951-8320(03)00058-9.

[41] Van Trees HL. Detection, estimation, and modulation theory. New York: Wiley; 2001. Sensor detection theory and false alarm rates, Chapter 2.

[42] of the Army D, editor. General Engineering. FM 3-34.400 / ATP 3-34.40 / MCWP 3-17.7. Department of the Army; 2014. Chapter 1: General Engineering in Full Spectrum Operations. Available from: https://www.globalsecurity.org/military/library/policy/army/fm/3-34-400/fm3-34-400.pdf.

[43] Paté-Cornell ME. Uncertainties in risk analysis: Six levels of treatment. Reliability Engineering & System Safety. 1996;54(2–3):95–111. https://doi.org/10.1016/S0951-8320(96)00067-1.

[44] Blake LV. Radar Range-Performance Analysis. Lexington, Mass: Lexington Books; 1980. Chapter 2: Detection probability and range.

[45] Richards MA. Fundamentals of radar signal processing. 2nd ed. New York: McGraw-Hill education; 2014.

[46] Bayes T. An Essay towards solving a Problem in the Doctrine of Chances. Philosophical Transactions of the Royal Society of London. 1763 Dec;53:370–418. https://doi.org/10.1098/rstl.1763.0053.

[47] Willis J, Blais CL, Burk RK, Gates BQ, Goerger NC, Keeter R, et al. Standards for the Mobility Common Operational Picture (M-COP): elements of ground vehicle maneuver; 2007. Available from: https://hdl.handle.net/11681/8519.

[48] Joseph A, Wiley A, Orr R, Schram B, Dawes J. The Impact of Load Carriage on Measures of Power and Agility in Tactical Occupations: A Critical Review. International Journal of Environmental Research and Public Health. 2018 Jan;15(1):88. https://doi.org/10.3390/ijerph15010088.

[49] Bonanno G. Axiomatic characterization of the AGM theory of belief revision in a temporal logic. Artificial Intelligence. 2007;171(2–3):144–160. https://doi.org/10.1016/j.artint.2006.12.001.

[50] Albus JS. 4-D/RCS reference model architecture for unmanned ground vehicles. In: Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065). vol. 4. San Francisco, CA, USA: IEEE; 2000. p. 3260–3265. Available from: http://ieeexplore.ieee.org/document/845165/.

[51] United States Air Force, LeMay Center for Doctrine Development and Education. Air Force Doctrine Publication 3-60: Targeting. United States Air Force; 2021. AFDP 3-60. Pages 19-21; last updated March 2019 version. Available from: https://www.doctrine.af.mil/Portals/61/documents/AFDP_3-60/3-60-AFDP-TARGETING.pdf.

[52] Endsley MR. Toward a theory of situation awareness in dynamic systems. Human Factors. 1995;37(1):32–64. https://doi.org/10.1518/001872095779049543.

[53] Shannon CE. A Mathematical Theory of Communication. Bell System Technical Journal. 1948;27(3):379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.

[54] Brown N, Sandholm T. Superhuman AI for multiplayer poker. Science. 2019;365(6456):885–890. Strategic adaptation in adversarial settings. https://doi.org/10.1126/science.aay2400.

[55] Washburn A, Kress M. Combat Modeling. International Series in Operations Research & Management Science. Springer; 2013. Pages 111–132.

[56] Simon HA. A Behavioral Model of Rational Choice. Quarterly Journal of Economics. 1955;69(1):99–118. https://doi.org/10.2307/1884852.

[57] Department of Defense. DoD Climate Adaptation Plan. Office of the Under Secretary of Defense; 2019. Environmental impacts on military operations, pages 21-37. Available from: https://www.sustainability.gov/pdfs/dod-2024-cap.pdf.

[58] Bruneau M, Chang SE, Eguchi RT, Lee GC, O'Rourke TD, Reinhorn AM, et al. A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities. Earthquake Spectra. 2003;19(4):733–752. https://doi.org/10.1193/1.1623497.

[59] : FM 3-0: Operations. Adversary movement rates and operational tempo, Chapter 2, Pages 2-4. Available from: https://irp.fas.org/doddir/army/fm3-0.pdf.

[60] Koblick DC. Sensor Processing for Ground Optical Signal Architecture for Contributing Space-Based SSA Sensor Data. In: Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS 2014). Maui, Hawai'i: AMOS Tech; 2014. p. 52. Available from: https://amostech.com/TechnicalPapers/2014/Sensor_Processing/KOBLICK.pdf.

[61] Beard RW, McLain TW. Small unmanned aircraft: Theory and practice. Princeton University Press; 2012. Chapter 7, Pages 120-142. Available from: https://press.princeton.edu/books/hardcover/9780691149219/small-unmanned-aircraft.

[62] Valavanis KP, Vachtsevanos GJ. Handbook of Unmanned Aerial Vehicles. Springer; 2015. Chapter 74, Pages 1813-1816.

[63] Chai Y, Wang Y. Optimal Control of Information Diffusion in Temporal Networks. IEEE Transactions on Network and Service Management. 2023;20(1):104–119. https://doi.org/10.1109/TNSM.2022.3209994.

[64] Yamín D, Medaglia AL, Akkinepally AP. Reliable Routing Strategies on Urban Transportation Networks. Transportation Science. 2024;58(2):377–393. https://doi.org/10.1287/trsc.2023.0013.

[65] Goldsmith SF, Aiken AS, Wilkerson DS. Measuring Empirical Computational Complexity. In: Proceedings of the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering. ESEC-FSE '07. Cavat, Croatia: ACM Press; 2007. p. 395–404.