

Chance-Constrained Belief-Space Maneuver Planning for Autonomous Collision Avoidance Under Uncertainty

Grace Ra Kim

Stanford University

Duncan Eddy and Mykel J. Kochenderfer

Stanford University

ABSTRACT

Increasing conjunction frequency in low Earth orbit places growing pressure on spacecraft operators to determine not only whether an encounter requires mitigation, but whether sufficient information is available to commit to a maneuver. Delaying action may provide additional tracking information that reduces state uncertainty and resolves an apparent conjunction. However, waiting also reduces the time and maneuver authority available if intervention ultimately becomes necessary, potentially increasing the required propellant expenditure. This work formulates this information–action tradeoff as a belief-space planning problem for conjunctions between a maneuverable spacecraft and an unmaneuverable secondary object. The planner represents the uncertain orbital states as Gaussian beliefs and uses a chance-constrained belief-space Monte Carlo tree search framework to reason over possible future tracking updates before time of closest approach (TCA). Because these observations are not known in advance, the resulting belief at TCA is uncertain, as different possible observation sequences may produce different terminal belief states. A terminal chance constraint limits the probability of reaching TCA with a belief whose probability of collision exceeds a prescribed threshold, allowing the planner to wait when future tracking is likely to resolve the conjunction while requiring intervention when continued deferral becomes too risky. We evaluate the approach on eight historical conjunctions from NASA’s Conjunction Assessment Risk Analysis dataset. By varying the secondary-object measurement quality and tracking cadence, we generate a total of 96 distinct evaluation scenarios. Across the evaluated conditions, the planner reaches TCA without maneuvering in approximately 40% of episodes while maintaining no terminal collision-risk violations. In contrast, fixed-time rule-based maneuver policies resolve more encounters without maneuvering when intervention is deferred closer to TCA, but at the expense of increasing terminal risk violations. The ability to delay maneuver decisions depends strongly on tracking quality and measurement cadence, ranging from 76% under accurate, frequent measurements to approximately 18%–20% under the poorest tracking conditions. For the same physical conjunction, changing only the expected tracking quality can shift the preferred decision from continued waiting to delayed maneuver commitment or immediate intervention. These results show that tracking quality and frequency are not only inputs to collision-risk estimation: they can determine when intervention becomes necessary. Explicitly valuing future information therefore provides a principled mechanism for deciding when an avoidance decision can safely remain open and when uncertainty itself justifies action.

1. INTRODUCTION

Low Earth orbit (LEO) is becoming increasingly congested, increasing the demands placed on the systems and operators responsible for conjunction assessment. In 2023, the United States Space Force 18th Space Defense Squadron (18 SDS) generated approximately 600,000 Conjunction Data Messages (CDMs) per day, three times the average daily rate reported in 2020 [1]. Behind this aggregate volume, however, each conjunction unfolds as a relatively short sequence of evolving risk assessments. As new tracking data become available, state estimates and covariances are refined, collision-risk estimates are updated, and operators must continually reassess whether a maneuver is warranted. At the same time, the information available for any one encounter can be sparse. NASA’s Conjunction Assessment Risk Analysis (CARA) program reports that conjunctions are generally first identified within seven days of the time of closest approach (TCA), yet most have fewer than ten CDMs available over that period [2]. Collision avoidance therefore presents a sequential decision problem in which operators must determine not only whether an encounter poses sufficient risk to warrant action, but also when the available information is sufficient to commit to a maneuver.

This timing question is particularly important for conjunctions between a maneuverable spacecraft and an unmaneuverable secondary object. Although coordination between active operators has become increasingly important in dense lower-LEO constellation shells [3], ESA reports that conjunctions at higher LEO altitudes remain dominated by debris [4]. More generally, an “active” classification does not necessarily imply maneuverability, as ESA classifies an object as active based on the existence of a communication link and explicitly notes that such an object may nevertheless be nonmaneuverable. When the secondary cannot coordinate an avoidance action or provide an authoritative maneuver plan, the primary spacecraft bears the burden of conjunction mitigation, while knowledge of the secondary is expected to improve through future surveillance observations. In this setting, choosing to wait is a deliberate action. Delaying a maneuver allows the planner to incorporate additional tracking information before committing to an avoidance action, but simultaneously reduces the remaining time and fuel efficiency and mitigation effectiveness available if intervention ultimately becomes necessary.

Operational practice already reflects this tradeoff. Conjunction decisions are commonly informed by the probability of collision, P_c , evaluated against prescribed action criteria. For example, NASA specifies a baseline mitigation threshold of $P_c > 10^{-4}$ at the mitigation commitment point [5]. Operators also differ in maneuver timing, with some favoring early intervention and others delaying action to refine the estimated encounter geometry and collision risk [6]. Delaying intervention may also avoid the operational impact of an unnecessary maneuver when an uncertain conjunction later resolves. Existing collision avoidance research has developed methods for collision-risk assessment and fuel-efficient maneuver design under state uncertainty [7, 8, 9]. More recently, chance constraints have been incorporated directly into spacecraft collision-avoidance optimization. Ryu et al. formulate fuel-optimal debris avoidance with a distributionally robust chance constraint on collision probability [10], while Arzelier et al. use sample-average approximation to minimize maneuver cost subject to a prescribed collision-probability threshold [11]. Chance-constrained control has also been applied to spacecraft trajectory planning under navigation, maneuver-execution, and dynamical uncertainty [12]. These approaches demonstrate how collision risk can be treated explicitly as a probabilistic constraint rather than only as a term in the maneuver objective. In these formulations, however, uncertainty primarily enters the optimization of the maneuver or trajectory itself. Future observations are not treated as part of the decision process, and minimal explicit reasoning is performed about the information that could be gained at future stages.

A separate line of recent work has begun to treat collision avoidance as a sequential decision problem. Bourriez et al. formulate autonomous collision avoidance as a POMDP with imperfect debris monitoring [13]. Kim et al. extend this decision-theoretic perspective to multi-spacecraft collision avoidance, using a semi-decentralized POMDP to account for intermittent information sharing under realistic communication constraints [3]. Kuhl et al. consider online maneuver planning against a nonmaneuverable object using an MDP [14] and Ferrara et al. explicitly examine the trade between delaying a maneuver for improved conjunction information and preserving the opportunity for effective intervention [15]. These works establish that maneuver timing and evolving information belong within the decision process itself. They do not, however, impose the terminal safety requirement considered in this work: a bound on the probability, over possible future observations and resulting belief evolution, that the collision-risk criterion will be violated at TCA.

While these prior works motivate a sequential decision formulation of collision avoidance, they do not explicitly constrain risk over the possible belief states induced by future tracking observations. Because those observations are not known in advance, different measurement sequences may produce different beliefs at TCA, potentially resolving the conjunction or reinforcing the need to maneuver. The relevant safety question is therefore not only whether the current probability of collision satisfies a prescribed threshold, but how likely future observations are to produce a terminal belief that violates that safety threshold.

In this work, we formulate autonomous collision avoidance as a finite-horizon, belief-space planning problem for a maneuverable primary spacecraft and an unmaneuverable secondary. The planner maintains Gaussian beliefs over both orbital states and uses a chance-constrained Monte Carlo tree search (MCTS) [16, 17] planner to reason over possible future tracking observations before TCA. Because these observations are not yet known, each candidate action induces a distribution over possible terminal beliefs. We impose a chance constraint on this distribution, requiring the probability of reaching TCA with a belief whose collision probability exceeds a prescribed threshold to remain below a specified risk level. This construction separates the collision-probability threshold that defines an unsafe terminal belief from the likelihood that future observations lead to such a belief. By applying the constraint at TCA, the planner can intelligently continue to wait when additional tracking is likely to resolve the conjunction and commit to a maneuver when continued deferral would no longer satisfy the prescribed terminal risk requirement.

There are three main contributions of this work. First, we formulate the maneuverable–non-maneuverable maneuver decision as a terminal chance-constrained belief-space planning problem, allowing future tracking information to influence the decision to wait or maneuver before that information is observed. Second, we evaluate the resulting planner on real NASA CARA conjunction scenarios and compare its decisions against rule-based and optimistic greedy offline-planning policies, characterizing the trade between safely deferring a maneuver or committing to a maneuver early. Third, we examine how this behavior changes with measurement accuracy and tracking cadence, identifying the transition between tracking regimes in which uncertainty materially changes the maneuver decision. Collectively, these results establish chance-constrained belief-space planning as a principled framework for autonomous collision avoidance that explicitly accounts for the decision value of future tracking information.

2. PROBLEM FORMULATION

We formulate spacecraft collision avoidance as a finite-horizon sequential decision problem under evolving state-estimate uncertainty. The central decision is whether an avoidance maneuver is necessary and, if so, whether to execute it immediately or wait for additional tracking that may resolve the conjunction without intervention. Acting earlier provides more time for a maneuver to modify the encounter geometry, but requires committing before additional tracking can determine whether intervention is necessary. Waiting allows future observations to refine the estimated state and uncertainty, but reduces the remaining time available for intervention. We represent this information–action tradeoff as a belief-space decision problem and impose the probability of collision P_c as the probabilistic safety quantity.

2.1 Conjunction Scenario and Scope

We consider a conjunction between one maneuverable active spacecraft, denoted by s , and one nonmaneuvering secondary resident space object, denoted by d . The secondary may represent orbital debris, a rocket body, or another resident space object for which no avoidance action is modeled. The Cartesian states of the primary and secondary objects in the Earth-centered inertial (ECI) frame are

$$x^s = [r_x^s \ r_y^s \ r_z^s \ v_x^s \ v_y^s \ v_z^s]^\top \in \mathbb{R}^6, \quad x^d = [r_x^d \ r_y^d \ r_z^d \ v_x^d \ v_y^d \ v_z^d]^\top \in \mathbb{R}^6 \quad (1)$$

where the first three components are position and the final three components are velocity. The planning horizon for a conjunction event starts with the generation of an initial CDM for the conjunction and terminates at the time of closest approach. We denote the time remaining to TCA by τ , such that $\tau = 0$ corresponds to the TCA.

This work considers conjunctions for which a two-dimensional encounter-plane probability of collision provides an appropriate representation of risk. In particular, we consider short-duration encounters for which the relative motion near TCA is sufficiently well represented by the assumptions underlying the 2D collision-probability model. Encounters outside this validity regime were not considered in the present experiments. The planning framework itself is not restricted to this regime, but extension to these encounters would require an appropriate alternative risk metric, such as a higher-fidelity P_c method, which may come at greater computational cost.

2.2 Belief-MDP Formulation

Sequential collision avoidance requires a planner to repeatedly decide to maneuver or wait as the estimated encounter advances towards TCA. If the physical state of the conjuncting objects were known exactly, this problem could be represented as a Markov decision process (MDP), in which each action is selected from the current state to optimize a long-term objective [18]. In practice, however, the primary spacecraft and secondary object orbital states are not known exactly. They are inferred from imperfect tracking measurements and are therefore uncertain.

This uncertainty makes the underlying problem partially observable. A partially observable Markov decision process (POMDP) extends the MDP framework by allowing the planner to receive noisy or incomplete observations of an underlying physical state rather than observing that state directly [19, 20]. The planner must therefore base its decisions on the belief, or the distribution of potential states.

A POMDP can be reformulated as an equivalent MDP over belief-states, where each belief is a probability distribution over the possible underlying physical states [20, 21]. This representation is particularly useful for collision avoidance because P_c depends on both the estimated encounter geometry and its associated uncertainty. As new observations are received, the belief is updated, allowing the planner to recompute P_c and reason directly over the evolving state

uncertainty. We therefore formulate collision avoidance as a finite-horizon belief-MDP $\mathcal{M}_b$, with the current belief b serving as the state for both risk evaluation and maneuver planning

$$\mathcal{M}_b = (\mathcal{B}, \mathcal{A}, T, \mathcal{R}, \mathcal{H}) \quad (2)$$

where $\mathcal{B}$ is the belief space, $\mathcal{A}$ is the action space, T is the transition model over beliefs, $\mathcal{R}$ defines the reward or decision objective, and $\mathcal{H}$ is the planning horizon. Decisions are made at a finite set of prescribed times before TCA. We index these decision epochs by k , with monotonically decreasing τ_k denoting the time remaining to TCA at epoch k . At each epoch, the planner maintains a belief b_k representing its current probabilistic estimate of the conjunction state. Because future observations are not known when an action is selected, the same current belief and action may lead to different successor beliefs. The transition model T therefore represents the distribution over possible successor beliefs induced by the future observation. Once a particular observation is realized, the corresponding posterior belief is obtained through a belief update model.

2.2.1 Belief Space

We represent the state uncertainty of the primary and secondary objects at decision epoch k by multivariate Gaussian beliefs, parameterized by their respective state means and covariances:

$$x_k^s \sim \mathcal{N}(\mu_k^s, \Sigma_k^s), \quad x_k^d \sim \mathcal{N}(\mu_k^d, \Sigma_k^d) \quad (3)$$

where $\mu_k^s, \mu_k^d \in \mathbb{R}^6$ are the estimated Cartesian states and $\Sigma_k^s, \Sigma_k^d \in \mathbb{R}^{6 \times 6}$ are the corresponding state covariance matrices. The belief $b_k \in \mathcal{B}$ used for planning is therefore represented by

$$b_k = (\mu_k^s, \Sigma_k^s, \mu_k^d, \Sigma_k^d, \tau_k) \quad (4)$$

which contains the estimated state and uncertainty of both objects together with the time remaining to TCA. We assume independent state estimation errors between the two objects, eliminating the need to maintain cross-correlation terms and reducing the dimensionality of the belief representation.

2.2.2 Action Space

At each decision epoch, the planner selects an action $a_k \in \mathcal{A} = \{\text{WAIT}, \text{MANEUVER}\}$, corresponding to either waiting for additional information or executing an avoidance maneuver. Here, WAIT applies no velocity change and MANEUVER applies a fixed impulsive maneuver to the active spacecraft. The maneuver magnitude is fixed at $\Delta v^{\text{man}} = 0.1$ m/s, a conservative but operationally plausible magnitude for a low-Earth-orbit collision-avoidance maneuver [9, 22], and is applied in the prograde direction according to

$$\Delta \mathbf{v}_s = \Delta v^{\text{man}} \frac{\mathbf{v}_s}{\|\mathbf{v}_s\|} \quad (5)$$

When the planner selects the MANEUVER action, this impulsive velocity change is applied immediately at the current decision epoch. The planner therefore determines, at each decision epoch, whether to execute the fixed maneuver or continue waiting, thereby selecting the maneuver timing from the discrete set of available decision epochs. Holding the maneuver magnitude and direction fixed isolates the effect of maneuver timing across the discrete decision epochs under evolving state uncertainty.

2.2.3 Transition and Observation Model

The belief evolves from one decision epoch to the next according to the selected action and the tracking information received before the next decision. Given the current belief b_k and action a_k , the belief transition model $T(b_{k+1} | b_k, a_k)$ describes the distribution over possible successor beliefs b_{k+1} at the next decision epoch. These transitions occur over the finite set of prescribed times $\mathcal{E} = \{\tau_0, \tau_1, \dots, \tau_{\text{TCA}} = 0\}$, where τ_k denotes the time remaining to TCA at epoch k . Intermediate decision epochs occur at a prescribed fixed cadence corresponding to scheduled secondary tracking

updates. At each decision epoch, the secondary belief is updated using the newly received measurement, the primary belief is updated using its assumed onboard navigation information, and the planner selects WAIT or MANEUVER.

Future observations are not known when a_k is selected, so the same current belief and action may lead to different successor beliefs depending on the observation received before the next decision epoch. At each scheduled tracking opportunity, the current Gaussian belief is propagated under action a_k to the next decision epoch and updated using the received observation z_{k+1} through a Kalman measurement update. The propagation and measurement-update models are detailed in Section 3.1. During planning, future observations are sampled from the observation model, producing different possible successor beliefs and allowing the planner to reason over how future tracking may change the estimated encounter geometry and uncertainty before TCA.

2.3 Chance-Constrained Collision Avoidance

The belief-MDP describes how the conjunction estimate may evolve under future actions and observations. We next define the collision-risk quantity associated with a belief and formulate maneuver selection subject to a chance constraint on the terminal collision risk.

2.3.1 Probability of Collision

When evaluating collision risk, the current belief b_k is propagated from decision epoch k to TCA to obtain the corresponding predicted state distribution at closest approach. For the short-duration conjunctions considered in this work, collision risk is evaluated in the two-dimensional encounter plane at TCA, defined as the plane normal to the relative velocity vector of the two objects at closest approach [23, 24]. Given a belief b_k at the current decision epoch, the primary and secondary state distributions are propagated to TCA and combined to form the relative-position distribution. This distribution is then projected onto the encounter plane. We denote the resulting encounter-plane relative position between the primary and secondary objects at TCA by the random variable $Y_k \in \mathbb{R}^2$, which is normally distributed as

$$Y_k \sim \mathcal{N}(\mu_{Y,k}, \Sigma_{Y,k}) \quad (6)$$

where $\mu_{Y,k}$ and $\Sigma_{Y,k}$ are the mean and covariance of the relative position between the two objects, projected onto the encounter plane at TCA. These quantities are obtained by propagating the belief at decision epoch k to TCA. Let ρ denote the combined hard-body radius of the two objects. The probability of collision associated with belief b_k is then

$$P_c(b_k) = \Pr(\|Y_k\| \leq \rho) \quad (7)$$

where $\|Y_k\| \leq \rho$ denotes the event that the relative position at TCA lies within the hard-body collision region [23, 24]. Because the distribution of Y_k is determined by the belief b_k , $P_c(b_k)$ depends on both the estimated encounter geometry and its associated uncertainty. The specific encounter-plane construction and numerical method used to evaluate $P_c(b_k)$ are described in the Appendix.

2.3.2 Chance-Constrained Objective

At decision epoch k , the observations that will be received over the remainder of the planning horizon are not yet known. Different observation sequences may produce different beliefs at TCA $\tau = 0$. We denote this uncertain future belief by the random variable $B_{\tau=0}$. Once a particular sequence of observations is realized, the resulting terminal belief is denoted by $b_{\tau=0}$. A realized terminal belief violates the prescribed collision-risk criterion when

$$P_c(b_{\tau=0}) > \delta \quad (8)$$

where δ is the maximum acceptable probability of collision. From the current decision epoch k , however, it is not yet known whether this event will occur because the terminal belief $b_{\tau=0}$ depends on future observations. The chance constraint therefore bounds the probability of the occurrence of this violation event with

$$\Pr(P_c(B_{\tau=0}) > \delta \mid b_k, a_k) < \alpha \quad (9)$$

where α specifies the maximum acceptable probability of violating the collision-risk criterion δ over future observations and the resulting belief evolution. The two parameters therefore operate at different levels: δ defines what

constitutes an unsafe terminal belief with probability of collision, while α defines the maximum probability the planner may tolerate of reaching an unsafe terminal belief. In this work, we prescribe the probability of collision threshold as $\delta = 10^{-5}$, which is consistent with the range of (P_c) criteria reported by LEO operators and equals the median high-interest threshold in the operator survey of Alfano et al. [6]. The threshold is also more conservative than NASA's (10^{-4}) baseline mitigation criterion [5]. The planner's chance-constraint risk level is set to $\alpha = 0.05$, balancing the desired confidence in constraint satisfaction with the need to retain feasible candidate actions. Because the chance constraint is imposed only on the terminal collision risk at TCA, intermediate beliefs along the planning horizon may have $P_c(b_k) > \delta$. These intermediate values do not by themselves make an action infeasible, since subsequent tracking observations may change the terminal belief and reduce the probability of violating the collision-risk criterion at TCA. Evaluating the constraint at TCA therefore allows the planner to account for future belief evolution without requiring a maneuver solely because an intermediate estimate exceeds the collision-risk threshold.

When candidate actions satisfy the terminal chance constraint, the planning objective is to balance expected terminal collision risk against maneuver expenditure. For a candidate action a_k selected from belief b_k , we define the root decision objective as

$$J(b_k, a_k) = -r_{\text{pc}} \mathbb{E}[P_c(B_{\tau=0}) \mid b_k, a_k] - r_{\text{man}} \mathbb{1}[a_k = \text{MANEUVER}] \quad (10)$$

where r_{pc} weights the expected terminal collision probability and r_{man} penalizes whether or not a maneuver was executed. The expectation $\mathbb{E}[P_c(B_{\tau=0}) \mid b_k, a_k]$ represents the average terminal collision probability over the possible terminal beliefs that may result after selecting action a_k from the current belief b_k . This quantity differs from the violation probability used by the chance constraint: the violation probability measures how often the terminal collision probability exceeds δ , whereas the expected terminal collision probability captures the average magnitude of collision risk across possible terminal outcomes.

3. METHODOLOGY

This section describes the planning and belief-propagation methodology used to solve the finite-horizon belief-MDP formulated in Section 2. We use a chance-constrained receding-horizon Monte Carlo tree search (MCTS) [16] planner to sample rollouts of possible future action-observation sequences from the current belief through TCA. The resulting terminal beliefs are used to estimate, for each candidate root action, both the probability of violating the prescribed collision-risk threshold and the expected terminal probability of collision. The former determines chance-constrained feasibility, while the latter is used together with maneuver cost to rank feasible actions. The overall planning architecture is illustrated in Figure 1.

3.1 Belief Propagation and Measurement Update

Each belief-space transition propagates the primary and secondary Gaussian beliefs to the next decision epoch and incorporates the available tracking information. For each object $j \in \{s, d\}$, the predicted mean and covariance are propagated according to

$$\mu_{k+1}^{j,-} = f^j(\mu_k^j, a_k, \Delta t_k), \quad \Sigma_{k+1}^{j,-} = \Phi_k^j \Sigma_k^j (\Phi_k^j)^\top \quad (11)$$

where f^j denotes the nonlinear orbital propagation model for object j , and Φ_k^j is the corresponding state-transition matrix over the interval between successive decision epochs. For the primary spacecraft, f^s applies the fixed prograde impulse when $a_k = \text{MANEUVER}$ before propagating the state. The secondary remains nonmaneuvering. At each tracking opportunity, object j is observed through the measurement model

$$z_{k+1}^j = h(x_{k+1}^j) + v_{k+1}^j, \quad v_{k+1}^j \sim \mathcal{N}(0, R^j) \quad (12)$$

where R^j is the measurement covariance. Observations are modeled as direct measurements of the full Cartesian state in ECI coordinates, such that

$$h(x^j) = x^j, \quad H^j = \frac{\partial h}{\partial x^j} = I_6 \quad (13)$$

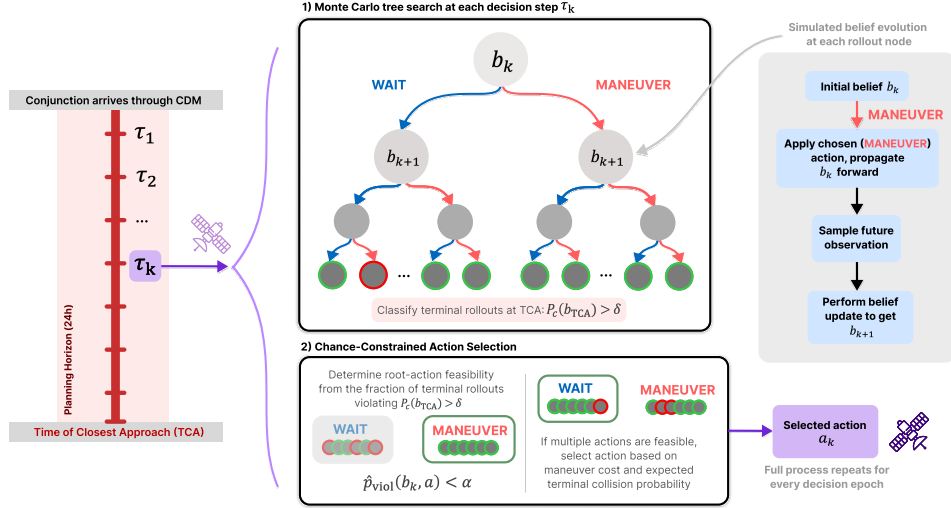


Fig. 1: Receding-horizon chance-constrained belief-space planning architecture. At each decision epoch, MCTS samples possible future belief evolution under the candidate root actions and evaluates the resulting terminal collision probabilities at TCA. The fraction of terminal rollouts exceeding δ determines chance-constrained feasibility under risk level α , after which feasible actions are ranked according to expected terminal collision risk and maneuver cost. Only the selected root action is executed before the belief is updated and planning is repeated.

and the measurement update therefore follows a simplified version of the extended Kalman filter equations

$$K_{k+1}^j = \Sigma_{k+1}^{j,-} \left[\Sigma_{k+1}^{j,-} + R^j \right]^{-1} \quad (14)$$

$$\mu_{k+1}^j = \mu_{k+1}^{j,-} + K_{k+1}^j \left(z_{k+1}^j - \mu_{k+1}^{j,-} \right) \quad (15)$$

$$\Sigma_{k+1}^j = \left(I_6 - K_{k+1}^j \right) \Sigma_{k+1}^{j,-} \quad (16)$$

where the simplification follows from $H^j = I_6$. The measurement model is linear, while the belief propagation follows the nonlinear orbital dynamics. During planning, future measurements are not known and are therefore sampled from the observation model. Each sampled measurement produces a different posterior belief, providing the observation-dependent branching used by the belief-space MCTS planner. The measurement-uncertainty models used to define R^j , along with the tracking schedules considered in the experiments, are described in Section 4.2.

Each terminal belief is propagated to TCA to evaluate $P_c(b_k)$, which is used for chance-constrained feasibility and root-action ranking. The encounter-plane construction and Elrod evaluation are described in the Appendix.

3.2 Belief-Space Monte Carlo Tree Search

The distribution of possible terminal beliefs $B_{\tau=0}$ cannot be enumerated directly because future tracking observations may produce different posterior beliefs over the remaining planning horizon. We therefore use Monte Carlo tree search to sample possible future action-observation sequences from the current belief through TCA. Starting from b_k , each MCTS rollout generates one possible sequence of future beliefs and terminates with a realization $b_{\tau=0}$. Repeated rollouts therefore provide additional samples of the possible terminal belief evolution used to evaluate the chance-constrained decision problem.

To guide tree-search exploration, we construct an MCTS reward inspired by the maneuver-cost and terminal collision-risk components of the root decision objective J in Equation (10). Within the tree, each rollout accumulates a fixed penalty r_{man} for each maneuver executed and, at TCA, a terminal collision-risk penalty $r_{\text{pc}} P_c(b_{\tau=0})$. Rollouts terminating above the collision-risk threshold additionally incur a large fixed penalty r_{viol} . The sample-average return

following action a from belief b defines the MCTS search value $Q(b, a)$. This value is then used to guide tree traversal according to the upper-confidence-bound criterion [25]

$$a_{\text{sel}} = \arg \max_{a \in \mathcal{A}} \left[Q(b, a) + c \sqrt{\frac{\ln N(b)}{N(b, a)}} \right] \quad (17)$$

where $N(b)$ is the number of visits to belief node b , $N(b, a)$ is the number of times action a has been selected from that node, and the constant c controls the exploration–exploitation tradeoff. Because future observations are continuous, progressive widening [26] limits the number of observation-dependent successor beliefs expanded from each action. We highlight that the search value $Q(b, a)$ is used only to guide sampling and exploration within the MCTS tree, whereas the terminal rollout outcomes are used to evaluate chance-constrained feasibility and final selection of feasible root actions according to $J(b_k, a_k)$, as described in the following section. As an ablation, we also consider direct root action selection using $Q(b_k, a)$ rather than J which we explore further in Section 5.1.

The MCTS planner is executed in a receding-horizon manner. At each decision epoch, it evaluates possible belief evolution through TCA and selects the root action a_k^* using the chance-constrained procedure in Section 3.3. Only this root action is executed. After propagating to the next decision epoch and incorporating the new tracking observation, the posterior belief b_{k+1} becomes the root of a new MCTS search over the shortened horizon. This process repeats until TCA, allowing the planner to revise whether to wait or maneuver as new information becomes available.

3.3 Chance-Constrained Root Action Selection

The terminal beliefs generated by MCTS are used to evaluate the chance constraint for each candidate root action. Let $\mathcal{L}(a_k)$ denote the set of rollouts beginning with action a_k , and let $L(a_k) = |\mathcal{L}(a_k)|$ denote the number of corresponding terminal outcomes. The violation probability introduced in Section 2.3.2 is estimated from the fraction of these rollouts whose terminal collision probability exceeds δ

$$\hat{p}_{\text{viol}}(b_k, a_k) = \frac{1}{L(a_k)} \sum_{\ell \in \mathcal{L}(a_k)} \mathbb{1} \left[P_c \left(b_{\tau=0}^{(\ell)} \right) > \delta \right] \quad (18)$$

where $b_{\tau=0}^{(\ell)}$ is the terminal belief produced by rollout ℓ . A candidate action is considered feasible when

$$\hat{p}_{\text{viol}}(b_k, a_k) < \alpha \quad (19)$$

under the given chance constraint. The same terminal outcomes are used to estimate the expected terminal collision probability appearing in the root decision objective from Equation (10)

$$\hat{\mathbb{E}}[P_c(B_{\tau=0}) \mid b_k, a_k] = \frac{1}{L(a_k)} \sum_{\ell \in \mathcal{L}(a_k)} P_c \left(b_{\tau=0}^{(\ell)} \right) \quad (20)$$

where these two estimates serve different purposes: $\hat{p}_{\text{viol}}$ determines whether an action satisfies the chance constraint, while the expected terminal collision probability characterizes the average terminal risk used to rank feasible actions. Substituting this estimate into the root decision objective in Equation (10) gives its Monte Carlo estimate $\hat{J}(b_k, a_k)$. The set of chance-constrained feasible actions is

$$\mathcal{A}_{\text{feas}}(b_k) = \{a \in \mathcal{A} \mid \hat{p}_{\text{viol}}(b_k, a) < \alpha\} \quad (21)$$

and the action executed at decision epoch k is selected according to

$$a_k^* = \arg \max_{a \in \mathcal{A}_{\text{feas}}(b_k)} \hat{J}(b_k, a) \quad (22)$$

and this construction allows for the separation of probabilistic safety from action ranking. The chance constraint first determines which candidate actions are admissible based on the sampled distribution of terminal collision risk. Afterwards, $\hat{J}$ selects among the feasible actions based on expected terminal risk and maneuver expenditure. However, if neither candidate action satisfies the chance constraint, the feasible set is empty. This may occur when uncertainty in the current belief prevents either action from satisfying the prescribed confidence level, even when acceptable

outcomes may exist. In this case, the planner instead selects the action with the lowest estimated expected terminal collision probability

$$a_k^* = \arg \min_{a \in \mathcal{A}} \hat{\mathbb{E}} [P_c(B_{\tau=0}) \mid b_k, a] \quad (23)$$

allowing the planner to select the least-risky action when neither available decision is predicted to satisfy the prescribed chance constraint under the current belief.

4. EXPERIMENTS

This section describes the conjunction scenarios, tracking conditions, and comparison policies used to evaluate the proposed chance-constrained planner. The NASA Conjunction Assessment Risk Analysis Tools dataset [27] provides the basis for the evaluation, with a subset of encounters selected to match the scope of this work. Measurement quality and tracking cadence are varied to examine how the information available to the planner affects its decisions, yielding 96 evaluation configurations (8 conjunctions $\times$ 3 measurement-quality levels $\times$ 4 cadences). Baseline policies and evaluation metrics are then introduced to characterize maneuver timing, collision risk, and propellant expenditure.

Orbital states are propagated numerically using the Brahe astrodynamics library [28]. Both objects are propagated in the Earth-centered inertial frame under a force model including a spherical-harmonic geopotential, atmospheric drag, third-body perturbations, and solar radiation pressure, with per-object ballistic parameters taken from the conjunction data message. Covariances are propagated by the linearized mapping obtained from the same numerical propagation. The same propagator is used to map each belief forward to TCA for collision-probability evaluation, and to back-propagate the CDM state and covariance from TCA to the start of the planning horizon.

4.1 Conjunction Dataset and Scenario Selection

We evaluate the proposed planner using conjunctions drawn from the NASA CARA Analysis Tools dataset [27]. The dataset contains historical conjunctions involving operational spacecraft between 2020 and 2023, provided as Conjunction Data Messages. Each CDM provides the primary and secondary state estimates and associated covariances at TCA together with encounter quantities including miss distance, relative velocity, combined hard-body radius, and a reference probability of collision.

The CDM state and covariance information is used to initialize the conjunction beliefs evaluated by the planner. Because the CDMs provide state estimates at TCA, the object states and covariances are propagated backward from TCA to the beginning of the planning horizon. Forward propagation of these beliefs then produces the evolving state-estimation uncertainty considered throughout the decision process while retaining the encounter geometry and covariance structure of the original conjunction.

As defined in Section 2.1, this work targets conjunctions for which the two-dimensional encounter-plane probability of collision provides an appropriate representation of collision risk. We apply a two-dimensional P_c usage-validity assessment, following NASA CARA’s implementation recommendations and usage boundaries [27], and exclude conjunctions outside the resulting validity regime. We additionally focus on conjunctions involving a nonmaneuvering debris secondary, consistent with the maneuverable-primary and nonmaneuvering-secondary problem considered in this work. The resulting evaluation set contains 8 debris-secondary conjunctions near the operational maneuver-decision window. Table 1 summarizes the selected encounters and their initial characteristics.

4.2 Tracking and Measurement Configurations

To evaluate how tracking information influences maneuver decisions, we independently vary the quality and cadence of measurements available for the secondary object. Measurement quality determines the uncertainty associated with each observation, while cadence determines how frequently new information becomes available before TCA.

Observations are represented in the filter as direct measurements of the full ECI state, with $H = I_6$. The measurement uncertainty, however, is constructed first in the secondary object’s local radial–transverse–normal (RTN) frame. This allows the state-space covariance to retain the range-versus-cross-range structure of the underlying radar measurements, with tighter uncertainty in the range direction and larger uncertainty in the two cross-range directions. Because the belief state is maintained in ECI, the resulting RTN covariance is rotated into ECI before being supplied to the extended Kalman filter. The three measurement-quality levels in Table 2 are derived from calibrated Space Surveillance Network (SSN) radar performance reported by Vallado [29] and accessed through brahe 1.7.0 [28]. The best, median,

Table 1: NASA CARA conjunction dataset selected for experimental evaluation.

Primary (NORAD)	Secondary (NORAD)	t_0 (h)	P_c^{CDM}	Miss (m)	HBR (m)
OCO-2 (40059)	IRIDIUM 33 DEB (35921)	21.8	7.9×10^{-4}	448	6.0
METOP-B (38771)	FENGYUN 1C DEB (30802)	25.1	1.6×10^{-3}	148	10.0
CALIPSO (29108)	COSMOS 2251 DEB (34995)	26.3	1.8×10^{-3}	197	14.8
NPP (37849)	THOR ABLESTAR DEB (13512)	26.5	1.1×10^{-2}	99	6.0
NOAA 19 (33591)	COSMOS 1275 DEB (42216)	27.0	4.5×10^{-3}	74	6.0
NOAA 18 (28654)	DMSP 5D-2 F12 DEB (41835)	27.3	5.0×10^{-3}	21	6.0
WORLDVIEW-3 (40115)	FENGYUN 1C DEB (30660)	27.7	1.1×10^{-4}	405	20.0
TERRA (25994)	CZ-4 DEB (26132)	59.1	1.2×10^{-3}	25	15.0

8 conjunctions; operational LEO payload primaries; debris secondaries; $\delta = 10^{-5}$

and worst configurations correspond to the 10th, 50th, and 90th percentiles of the radar noise distribution, respectively, and represent favorable, typical, and degraded tracking conditions rather than individual sensors.

For each quality level, range uncertainty defines the radial position uncertainty, while angular uncertainty is mapped to a representative cross-range uncertainty according to

$$\sigma_{\text{cross}} = \sigma_{\text{ang}} d \quad (24)$$

using a representative sensor-target slant range of $d = 1200$ km, approximately corresponding to a 20° elevation angle for a target near 550 km altitude. The measurement covariance is then constructed as

$$R_{\text{RTN}} = \text{diag}(\sigma_{\text{range}}^2, \sigma_{\text{cross}}^2, \sigma_{\text{cross}}^2, \sigma_v^2, \sigma_v^2, \sigma_v^2) \quad (25)$$

where $\sigma_v = 0.1$ m/s and the radar range standard deviation σ_{range} is used to define the radial position uncertainty. The covariance is rotated into ECI using the secondary object’s reference state at the CDM epoch and held fixed for subsequent measurement updates.

For each measurement-quality level, we additionally vary the secondary tracking cadence over 2, 4, 8, 24 h. These cadences represent idealized tracking schedules in which secondary measurements become available at prescribed intervals rather than through explicit sensor visibility or tasking models. The scheduled measurement times define the planner’s decision epochs, such that each new secondary measurement triggers a belief update followed by a new maneuver decision. Shorter cadences therefore provide more frequent opportunities to update the conjunction belief and reconsider whether to maneuver, while longer cadences require the planner to wait longer between decisions. The primary spacecraft is assumed to have access to its own GPS-grade navigation solution (10 m, 1σ) [30], which is updated at each of these decision epochs. Together, the measurement-quality and cadence sweeps allow us to evaluate when additional tracking information is sufficiently informative to justify delaying an avoidance maneuver.

Table 2: Secondary-object measurement-quality configurations derived from representative SSN radars [29]

Quality	Percentile	SSN radar noise		Mapped state σ (RTN)	
		Range (m)	Mean az/el (deg)	Radial (m)	Transverse, normal (m)
Best	10th	26.0	0.0115	26.0	241
Median	50th	50.0	0.0224	50.0	469
Worst	90th	140.3	0.0477	140.3	999

4.3 Baseline Policies and Ablations

We compare the proposed chance-constrained MCTS planner against baseline decision policies designed to isolate the benefit of reasoning over future belief evolution when determining maneuver timing. Each baseline is evaluated on the

same conjunctions, with the same measurement schedule, the same observation noise realizations, and the same belief update as the planner, so that any difference in outcome is attributable to the decision rule alone.

The *rule-based* family holds choosing to maneuver until the time to TCA falls below a threshold τ_{act} , at which point it maneuvers only if the current belief collision probability exceeds δ ; it then executes a single burn. This represents the common operational pattern of screening against a risk threshold on a fixed schedule, and reads P_c from the same beliefs available to the chance-constrained planner. We consider decision times of 28, 12, 6, and 3 h before TCA, spanning early intervention through increasingly delayed responses. Sweeping τ_{act} traces the trade-off between acting early on poor information and deferring past the window in which a maneuver remains effective.

The *optimistic greedy offline plan* represents a non-adaptive policy that commits to a maneuver strategy under favorable assumptions about future tracking. At the initial decision epoch, the policy propagates the WAIT trajectory to TCA while applying measurement updates according to the prescribed tracking schedule of 2, 4, 8, or 24 h. Each future measurement is set equal to the dynamics-predicted state, modeling an ideal measurement with no noise, so it contracts the covariance without updating the predicted trajectory. This represents planning under the optimistic, idealized assumption that the dynamics prediction is perfectly accurate, and future measurements will perfectly align, simply reducing uncertainty as TCA approaches. The policy then commits to a single decision under this optimistic forecast: if waiting is predicted to achieve $P_c < \delta$ at TCA, it commits to WAIT for the entire episode; otherwise it maneuvers immediately. This policy decision is made *a priori* and does not change during evaluation.

This baseline is intentionally optimistic about the information that future tracking will provide. It assumes that subsequent measurements will reduce uncertainty without shifting the trajectory prediction (mean state). During execution, however, measurements are drawn from the stochastic observation model and may not continue to confirm the initial trajectory prediction. A plan that appears safe under the optimistic forecast may therefore fail to achieve the anticipated risk reduction when the realized measurements are less favorable. The baseline consequently tests the danger of committing too strongly to an expected improvement in state knowledge. By pushing deferral as far as the favorable measurement forecast permits, it provides a deliberately aggressive reference for how many maneuvers could be avoided if future tracking perfectly supported a given prediction. Its failures illustrate the corresponding cost of over-confidence in a particular, potentially incorrect, predicted future evolution of a conjunction: a plan judged sufficiently mitigated in advance may no longer satisfy the collision-risk threshold once the actual measurements are incorporated.

Finally, as an ablation, we apply the *soft P_c penalty* of the chance-constrained action-selection rule, where we evaluate the same MCTS planner selecting the root action by $\arg \max_a Q(b_k, a)$. Here the collision constraint enters only softly, as the penalty r_{viol} embedded in the backed-up value, rather than as an explicitly estimated terminal violation probability enforced as a feasibility requirement. Comparing the two isolates the effect of separating chance-constrained feasibility from the search value used to guide tree exploration. Parameters held fixed across the test sweep are summarized in Table 3. All policies use the same fixed prograde maneuver magnitude and collision-risk thresholds.

Table 3: Fixed planner configurations.

Parameter	Value
Maneuver magnitude (Δv)	0.1 m s^{-1}
Collision-risk threshold (δ)	10^{-5}
Chance-constraint level (α)	0.05
Collision-risk weight (r_{pc})	10^6
Terminal violation penalty (r_{viol})	10^4
Maneuver cost (r_{man})	10
MCTS exploration constant (c)	10
MCTS rollouts per decision	100
Full Episode Seed Runs	5

5. RESULTS

The results evaluate three aspects of the proposed planner: its ability to safely defer maneuvers relative to the baseline policies, the sensitivity of this behavior to tracking quality, and the mechanism by which future information influences the maneuver decision. Results are aggregated across the CARA conjunctions, measurement realizations, and tracking configurations described in Section 4.

5.1 Safe Resolution Without Maneuver

We first examine how each policy balances resolving encounters without a maneuver against maintaining the prescribed collision-risk criterion at TCA. The results are shown in Figure 2.

The rule-based policies show how this balance changes with decision timing. The 28 h rule is the most conservative, producing no observed terminal violations but resolving relatively few encounters without a maneuver. As τ_{act} moves closer to TCA, the policy benefits from additional tracking information and therefore classifies more encounters as safe to leave unmitigated. However, the decision is fixed at τ_{act} and is not revised as later measurements arrive. Those measurements can subsequently shift the estimated encounter geometry and collision risk, causing some encounters that appeared safe at the decision time to terminate above δ . Thus, later decision times increase the number of scenarios resolved without a maneuver, but at the cost of higher terminal violation rates. For the 3 h rule, approximately 61–77% of encounters are resolved without maneuvering, with 10.6–11.2% of all episodes terminating above δ .

The optimistic greedy offline plan approaches the same decision from a different perspective. Rather than deferring the decision to a prescribed time, it commits at the initial epoch using an optimistic forecast of future tracking. Under this policy, 41–64% of encounters are ultimately resolved without a maneuver. When the anticipated measurements are subsequently realized, however, 0.6–4.5% of episodes terminate above δ . This result illustrates the consequence of relying too strongly on anticipated improvements in state knowledge. Future measurements can reduce uncertainty, but their realized innovations can also change the estimated encounter geometry and resulting collision risk.

The two belief-space planners occupy a different region of the comparison. Both maintain zero observed terminal violations across all three sensor-quality configurations, while still allowing a substantial fraction of encounters to reach TCA and resolve the conjunction without a maneuver. Within this zero-violation region, the chance-constrained




Fig. 2: Encounters resolved without a maneuver and P_c threshold violations at TCA for the proposed planner, its soft-penalty ablation, and comparison policies, separated by secondary measurement quality. Rule-based policies resolve more encounters without maneuvering as τ_{act} approaches TCA, but more frequently violate the chance constraint, violating 10.6–11.2% of episodes at the 3 h rule. The optimistic greedy offline plan commits to a single up-front decision under a forecast in which future measurements contract the covariance without shifting the state estimate; because that forecast overstates what tracking delivers, it still leaves 0.6–4.5% of episodes in violation once measurements are realized. The chance-constrained planner and its soft-penalty ablation both hold violations at zero, with the chance constraint resolving more encounters without a maneuver.


planner resolves more encounters without maneuvering than the soft- P_c penalty ablation. The chance constraint therefore improves the use of available tracking information without obtaining that improvement through an increase in observed terminal risk.

These results also clarify why the fraction of scenarios resolved without a maneuver should not be maximized in isolation. The 3 h rule-based policy resolves the largest share of encounters without maneuvering of any policy considered—76.9% under the best tracking, against 58.6% for the chance-constrained planner. However, this baseline does so while leaving 10.6% of episodes above δ at TCA, while the chance-constrained planners leave none. Some conjunctions remain above the prescribed risk criterion despite additional tracking and appropriately require a maneuver; a policy that avoids maneuvering in those cases has not resolved them, it has merely declined to act. The desired outcome is therefore not 100% no-maneuver resolution, but to identify which encounters can safely continue without intervention while retaining the ability to mitigate those that cannot. Figure 2 shows that the chance-constrained planner moves this balance toward greater no-maneuver resolution while satisfying the terminal chance constraint.

5.2 Effect of Tracking Quality and Frequency

The value of waiting depends strongly on both the quality of future measurements and how frequently they become available. Figure 3 compares the fraction of encounters resolved without a maneuver across sensor quality and measurement cadence. Because avoiding a maneuver is beneficial only when the resulting encounter remains safe, the figure reports this outcome together with the frequency of terminal collision-risk threshold violations.

For the chance-constrained planner, accurate and frequent measurements create substantial opportunity to defer intervention safely. Under the best sensor-quality configuration, 76% of encounters are resolved without a maneuver at a 2 h cadence and 72% at a 4 h cadence. This fraction decreases to 48% and 38% as the measurement cadence increases to 8 and 24 h, respectively. The same trend appears as measurement quality degrades. At median sensor quality, the fraction resolved without maneuvering falls from 60% at a 2 h cadence to 20% at 24 h. Under the worst sensor quality, the fraction resolved without maneuvering falls from 35% at a 2 h cadence to 20% at 24 h.



Cost of imperfect state knowledge, by measurement timing and sensor quality
Color = % of encounters resolved without a maneuver. White hatch = how often that left the operator P_c above threshold at TCA.

Fig. 3: Encounters resolved without a maneuver across secondary measurement cadence and sensor quality. Cell color gives the percentage resolved without maneuvering, and white hatching indicates terminal $P_c > \delta$ violations. Both belief-space planners maintain zero observed violations, while the 3 h rule-based policy and optimistic greedy offline plan reach violation rates of up to 17.5% and 14.3%, respectively. The chance-constrained planner outperforms the soft- P_c penalty ablation in eleven of twelve settings, with the largest gains under sparse measurements.

no-maneuver resolution remains at or below 35% across all cadences. Importantly, all conjunctions resolved without maneuvering have satisfied the terminal chance constraint. The interaction between measurement quality and cadence further shows that measurement frequency is most valuable when each observation is sufficiently informative. More frequent measurements substantially increase safe deferral under the best and median sensor qualities, whereas the benefit is smaller under the worst-quality configuration; even best-quality measurements at a 24 h cadence slightly outperform worst-quality measurements received every 2 h.

The comparison with the soft- P_c penalty ablation shows that explicitly constraining risk becomes particularly important when measurements are sparse. The chance-constrained planner resolves more encounters without maneuvering across the majority of tracking scenarios, with the largest separation occurring at the 8 and 24 h cadences. As measurements become more frequent, the two belief-space planners converge, indicating that frequent information can partially compensate for differences in how risk is incorporated into the decision rule.

The non-adaptive policies illustrate why no-maneuver resolution cannot be interpreted independently of safety. In particular, the optimistic greedy offline plan assumes favorable future measurement outcomes and commits at the initial decision epoch to the least-conservative plan that appears safe under that forecast. This assumption allows the plan to defer aggressively in several tracking regimes, but the apparent improvement comes with terminal threshold violations of up to 14.3%. The 3 h rule-based policy exhibits the same tradeoff more strongly, reaching violation rates of up to 17.5%. Thus, greater willingness to wait does not by itself indicate better use of future information.

Together, these results show that future tracking has substantial decision value, but that value depends on accounting for uncertainty in what future measurements will reveal. Accurate and frequent measurements allow the chance-constrained planner to avoid unnecessary maneuvers while maintaining the prescribed risk criterion. Conversely, planning as though future measurements will be sufficiently favorable can make deferral appear more attractive than it is during execution. The benefit of belief-space planning is therefore not simply that it waits for more information, but that it determines when waiting remains justified under uncertainty in that information.




Fig. 4: Belief collision probability through TCA for three representative conjunctions under best, median, and worst secondary measurement quality. Solid curves show the executed MCTS trajectory, while dashed curves show the counterfactual continued-WAIT trajectory. The examples illustrate encounters that resolve through tracking alone (left), require mitigation (center), and produce measurement-dependent maneuver decisions (right). These cases are drawn from across the sweep to illustrate the wait-and-measure behavior of the tree-based planner rather than to compare the chance-constrained and soft- P_c root-selection rules.

5.3 Information-Dependent Maneuver Timing

For the MCTS-based planners, the same conjunction can lead to different maneuver decisions as the quality of tracking information available before TCA changes. To examine this behavior in more detail, Figure 4 shows complete episodes for three representative conjunctions selected from across the sweep. The cases illustrate three distinct responses to additional tracking information: an encounter that resolves without a maneuver, one that requires mitigation in all tracking scenarios, and one for which the maneuver decision changes with measurement quality. The examples are intended to illustrate the common wait-and-measure behavior of the MCTS-based planners rather than highlight differences between the chance-constrained and soft- P_c root-selection rules.

For NORAD 40059 vs. 35921, continued tracking reduces the estimated collision probability below δ under all three measurement-quality configurations. The planner therefore reaches TCA without maneuvering, although the threshold is crossed earlier when measurements are more accurate. In contrast, NORAD 28654 vs. 41835 remains hazardous under the continued-WAIT trajectory for every tracking configuration, and the planner maneuvers at the initial decision epoch. NORAD 38771 vs. 30802 lies between these cases. With the best tracking configuration, the planner continues to wait and ultimately reaches TCA below δ without a maneuver. With median-quality tracking, it initially waits for additional information but later executes a maneuver as the belief evolves. Under the worst tracking configuration, it maneuvers immediately. The same conjunction therefore transitions from no maneuver, to a later maneuver, to immediate intervention as measurement quality decreases.

These examples show how tracking quality changes both whether and when a maneuver is required. For some conjunctions, additional measurements consistently provide enough information to continue without intervention, while others require mitigation regardless of tracking quality. Between these cases are conjunctions for which the maneuver decision depends directly on the information obtained as the belief evolves toward TCA.

5.4 Effect of Chance-Constrained Root Selection

The chance constraint and soft- P_c penalty use the same MCTS search and differ only in how they select the final action. The soft penalty averages the penalized returns across rollouts, so even a single threshold violation can strongly reduce an action’s value. The chance constraint instead considers how often violations occur, allowing the action to remain feasible as long as the estimated violation frequency remains below the chance constraint set by α .

Figure 5 illustrates this mechanism for a single root decision under best-quality tracking. WAIT violates the terminal threshold in one of its 27 rollouts, giving $\hat{p}_{\text{viol}} = 0.037 < \alpha = 0.05$, and therefore remains feasible. Its expected collision probability is $\mathbb{E}[P_c] = 4.9 \times 10^{-7}$, approximately twenty times below δ , so the chance-constrained rule selects WAIT as the lower-cost feasible action. Under the soft- P_c penalty, the same violating rollout drives $Q(\text{WAIT}) = -359$ compared with $Q(\text{MANEUVER}) = -13$, causing the planner to maneuver.

Across the 96 paired comparisons in the sweeps examined in Section 5.1 and Section 5.2, the two root rules disagree in 23 cases. In every disagreement, the chance-constrained rule selects WAIT while the soft- P_c penalty selects MANEUVER; the reverse is never observed. These disagreements concentrate near the constraint boundary, where a small number of violating rollouts can substantially reduce the mean return without causing the estimated violation frequency to exceed α . Despite selecting WAIT more often in these cases, the chance-constrained planner produces no observed increase in terminal threshold violations, with both belief-space planners maintaining zero violations across all three tracking tiers.

This behavior does not persist when the available tracking information no longer supports waiting. For the same conjunction under the worst sensor quality, WAIT has $\mathbb{E}[P_c] = 1.1 \times 10^{-3}$, more than two orders of magnitude above δ , and both rules select MANEUVER. The chance-constrained rule eliminates WAIT as infeasible, while the soft- P_c rule rejects it through its mean return. Thus, the chance-constrained root rule preserves WAIT when its estimated violation frequency remains within the prescribed bound, but selects a maneuver when the available tracking information no longer supports WAIT as a feasible action.


6. CONCLUSIONS

Collision avoidance decisions depend not only on the collision risk estimated from the information available now, but also on the information expected to become available before TCA. For encounters with an unmaneuverable secondary, this creates a fundamental tradeoff: maneuvering early preserves mitigation authority but commits to an action using a less-informed estimate, while waiting can resolve uncertainty but leaves less time to respond. The chance-constrained belief-space formulation developed here makes this tradeoff explicit by evaluating whether future tracking is sufficiently likely to produce an acceptable terminal belief before allowing the planner to defer intervention.


The experiments show that this distinction can materially change maneuver decisions. Across the evaluated NASA CARA conjunctions and tracking conditions, the chance-constrained planner resolved approximately 40% of episodes without a maneuver while maintaining zero terminal collision-risk violations. Rule-based policies deferred intervention more aggressively, but only by moving along a clear safety–deferral tradeoff, with threshold violation rates increasing to approximately 11% for the latest decision policy. The belief-space planner instead reasoned explicitly

NORAD 38771 vs 30802, Best Sensor Quality

Chance constraint: keep actions with $\hat{p}_{\text{viol}} < \alpha = 0.05$, then rank survivors by $E[P_c]$ and fuel



Soft penalty (ablation): no gate — take arg max Q , which carries the in-tree -10^4 penalty



● Feasible ($P_c < 10^{-5}$) ○ Transient warning ✗ Violating rollout at TCA (ringed) — WAIT branch — MANEUVER branch

Node size and edge width scale with MCTS visit count; deeper levels are imagined look-ahead rollouts, not committed actions.

Fig. 5: One MCTS root decision for NORAD 38771 vs. 30802 under best-quality tracking and an 8 h measurement cadence. Both root rules read the same 325-node search tree. The chance constraint accepts WAIT because only one of 27 rollouts violates the terminal threshold, giving $\hat{p}_{\text{viol}} = 0.037 < \alpha = 0.05$, and selects the lower-cost feasible action. The soft- P_c penalty instead allows this single violating rollout to dominate WAIT’s mean return and selects MANEUVER. For this conjunction, **waiting is ultimately preferable**: under high-quality tracking, subsequent measurements resolve sufficient uncertainty for the encounter to reach TCA below the collision-risk threshold without a maneuver. The chance constraint preserves this option despite a single adverse rollout that would otherwise trigger unnecessary intervention.

over possible future tracking outcomes to determine when deferral remained admissible under the prescribed terminal risk requirement.

More importantly, the results show that the value of waiting is determined by the tracking environment. Under the most accurate and frequent secondary tracking considered, 76% of episodes were resolved without maneuvering. As measurements became less accurate or less frequent, safe deferral decreased substantially, reaching approximately 18–20% under the least informative measurement conditions. This transition is not simply a change in the precision of the resulting collision-probability estimate. For the same NORAD 38771 and 30802 conjunction, the planner waited through TCA under high-quality tracking but maneuvered at the first decision epoch under low-quality tracking. The physical encounter was unchanged; only the information expected to arrive before TCA differed.

This result has a broader implication for autonomous conjunction management. Tracking performance and maneuver planning should not be treated as independent components of the response pipeline. The quality and timing of future observations determine how long an avoidance decision can safely remain open, and therefore have direct operational value in units of maneuver deferral. Conversely, degraded state knowledge carries a decision cost: when future observations cannot be expected to resolve the encounter with sufficient confidence, the planner must surrender the option to wait and commit earlier. The terminal chance constraint provides a principled boundary between these regimes by identifying when the uncertainty in future belief evolution makes continued deferral itself inadmissible.

The formulation developed here isolates this information–action tradeoff using a fixed impulsive maneuver, prescribed tracking opportunities, and short-duration conjunctions within the validity regime of the two-dimensional encounter-plane collision model. These assumptions provide a controlled setting in which the decision value of future information can be separated from continuous maneuver optimization and sensor scheduling. Extending the action space to continuous maneuver design and allowing future observations to be actively tasked would remove that separation and create a more general problem: jointly deciding when to observe, when to wait, and when and how to maneuver. In that setting, tracking resources and propellant become coupled resources for managing collision risk. The results presented here suggest that this coupling is consequential: improving knowledge of an encounter can eliminate the need to maneuver, while insufficient knowledge can itself become a reason to act.

ACKNOWLEDGMENTS

This research was supported by the Fannie and John Hertz Foundation and the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program. Specifically, this material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-25-C-B010 in the amount of currently negotiated tuition and stipend rates.

REFERENCES

- [1] Christian Ramos, Diana McKissock, Nicholas Herzer, Katelyn Eppich, Bailey Bourque, Joshua Parris, and Amos Alexander. Lessons Learned on Mega-Constellation Deployments and Impact on Space Domain Awareness. In *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, 2023.
- [2] Alinda K. Mashiku, Lauri K. Newman, and Dolan E. Highsmith. NASA Conjunction Assessment Risk Analysis (CARA) Compendium for Artificial Intelligence and Machine Learning for Satellite Collision Avoidance. In *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Wailea, HI, 2025.
- [3] Grace Ra Kim, Mahdi Al-Husseini, Duncan Eddy, and Mykel J. Kochenderfer. Semi-Decentralized Multi-Spacecraft Collision Avoidance under Communication Constraints. In *AAS/AIAA Astrodynamics Specialist Conference*, 2026.
- [4] ESA Space Debris Office. ESA’s Annual Space Environment Report. Issue 10.0 GEN-DB-LOG-00288-OPS-SD, European Space Agency, Darmstadt, Germany, May 2026.
- [5] National Aeronautics and Space Administration. NASA Spacecraft Conjunction Analysis and Collision Avoidance for Space Environment Protection. Technical Report NPR 8079.1, National Aeronautics and Space Administration, 2023.
- [6] Salvatore Alfano, Daniel L. Oltrogge, and Lorenzo Arona. Operators’ Requirements for SSA Services. *The Journal of the Astronautical Sciences*, 69:1441–1476, 2022.

- [7] Maruthi R. Akella and Kyle T. Alfriend. Probability of Collision Between Space Objects. *Journal of Guidance, Control, and Dynamics*, 23(5):769–772, 2000.
- [8] Kyle T. Alfriend, Maruthi R. Akella, Joseph Frisbee, James L. Foster, Deok-Jin Lee, and Matthew P. Wilkins. Probability of Collision Error Analysis. *Space Debris*, 1(1):21–35, 1999.
- [9] Claudio Bombardelli and Javier Hernando-Ayuso. Optimal Impulsive Collision Avoidance in Low Earth Orbit. *Journal of Guidance, Control, and Dynamics*, 38(2):217–225, 2015.
- [10] Kanghyun Ryu, Jean-Baptiste Bouvier, Shazaib Lalani, Siegfried Eggel, and Negar Zahedi Mehr. Risk-Sensitive Orbital Debris Collision Avoidance Using Distributionally Robust Chance Constraints. In *AIAA SCITECH 2025 Forum*, 2025.
- [11] Denis Arzelier, Fabrizio Dabbene, Mioara Joldeş, Martina Mammarella, Matthieu Masson, and Pema Mercereau-Boland. An Evaluation of Sample Average Approximation Applied to the Design of Impulsive Thrust Space Collision Avoidance Maneuvers. *IFAC-PapersOnLine*, 59(31):97–102, 2025.
- [12] Kenshiro Oguri. Chance-Constrained Control for Safe Spacecraft Autonomy: Convex Programming Approach. In *2024 American Control Conference (ACC)*, 2024.
- [13] Nicolas Bourriez, Adrien Loizeau, and Adam F. Abdin. Spacecraft Autonomous Decision-Planning for Collision Avoidance: A Reinforcement Learning Approach. In *International Astronautical Congress (IAC)*, 2023.
- [14] William Kuhl, Jun Wang, Duncan Eddy, and Mykel J. Kochenderfer. Markov Decision Processes for Satellite Maneuver Planning and Collision Avoidance. In *IEEE Aerospace Conference*, 2025.
- [15] Francesca Ferrara, Lander W. Schillinger Arana, Florian Dörfler, and Sarah H. Q. Li. A Markov Decision Process Framework for Early Maneuver Decisions in Satellite Collision Avoidance. In *AAS/AIAA Astrodynamics Specialist Conference*, 2025.
- [16] Cameron B. Browne, Edward Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.
- [17] Robert J. Moss, Arec Jamgochian, Johannes Fischer, Anthony Corso, and Mykel J. Kochenderfer. ConstrainedZero: Chance-Constrained POMDP Planning Using Learned Probabilistic Failure Surrogates and Adaptive Safety Constraints. In *International Joint Conference on Artificial Intelligence*, pages 6752–6760, 2024.
- [18] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York, NY, 1994.
- [19] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence*, 101(1–2):99–134, 1998.
- [20] Mykel J. Kochenderfer, Tim A. Wheeler, and Kyle H. Wray. *Algorithms for Decision Making*. The MIT Press, 2022.
- [21] Richard D. Smallwood and Edward J. Sondik. The Optimal Control of Partially Observable Markov Processes over a Finite Horizon. *Operations Research*, 21(5):1071–1088, 1973.
- [22] H. Klinkrad, J. R. Alarcon, and N. Sanchez. Collision Avoidance for Operational ESA Satellites. In *European Conference on Space Debris*, volume 587, page 509, 2005.
- [23] James Lee Foster and Herbert S. Estes. *A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles*. NASA, National Aeronautics and Space Administration, Lyndon B. Johnson Space Center, 1992.
- [24] Salvatore Alfano. A Numerical Implementation of Spherical Object Collision Probability. *The Journal of the Astronautical Sciences*, 53(1):103–109, 2005.
- [25] Levente Kocsis and Csaba Szepesvári. Bandit-Based Monte-Carlo Planning. In *European Conference on Machine Learning*, 2006.
- [26] Zachary Sunberg and Mykel Kochenderfer. Online Algorithms for POMDPs with Continuous State, Action, and Observation Spaces. In *International Conference on Automated Planning and Scheduling*, volume 28, pages 259–263, 2018.
- [27] NASA Conjunction Assessment Risk Analysis (CARA) Team. CARA Analysis Tools, 2026. Accessed: 2026-08-14.
- [28] Duncan Eddy and Mykel J. Kochenderfer. Brahe: A Modern Astrodynamics Library for Research and Engineering Applications. 2026.
- [29] David A. Vallado. *Fundamentals of Astrodynamics and Applications*. Space Technology Library. Microcosm Press, Hawthorne, CA, 4 edition, 2013.
- [30] André Hauschild and Oliver Montenbruck. Precise Real-Time Navigation of LEO Satellites Using GNSS Broad-

cast Ephemerides. *NAVIGATION*, 68(2):419–432, 2021.

- [31] Chris Elrod, James A. Stamey, and Matthew D. Hejduk. Computational Bayesian Methods Applied to Complex Problems in Bio and Astro Statistics. Technical report, 2019.

APPENDIX

Collision-Probability Evaluation. For each belief encountered during planning, the probability of collision is evaluated at TCA and used both to test chance-constrained feasibility and to score the decision objective. Starting from a belief b_k , the mean state and covariance of each object are propagated to TCA using the dynamics and covariance propagation of Section 3.1. Importantly, $\Sigma_{\tau=0}$ is obtained by propagating the *accumulated* belief covariance forward from b_k , so that the information gained from observations taken before t_k is reflected in the terminal risk; it is not reset to the initial covariance. The position covariances are combined to give the relative-position covariance

$$\Sigma_{\text{rel}} = \left(\Sigma_{\tau=0}^s\right)^{(r)} + \left(\Sigma_{\tau=0}^d\right)^{(r)} \quad (26)$$

where $\left(\Sigma_{\tau=0}^s\right)^{(r)}$ and $\left(\Sigma_{\tau=0}^d\right)^{(r)}$ are the position submatrices of the propagated primary and secondary covariances, respectively. The corresponding relative-position mean is

$$\mu_{\text{rel}} = \left(\mu_{\tau=0}^s\right)^{(r)} - \left(\mu_{\tau=0}^d\right)^{(r)} \quad (27)$$

Because the conjunctions considered here are short-duration, high relative-velocity encounters, the relative motion is approximately linear and the covariance approximately static over the encounter, permitting the standard two-dimensional reduction [23, 24]. Let

$$v_{\text{rel}} = \left(\mu_{\tau=0}^s\right)^{(v)} - \left(\mu_{\tau=0}^d\right)^{(v)}, \quad \hat{v}_{\text{rel}} = \frac{v_{\text{rel}}}{\|v_{\text{rel}}\|} \quad (28)$$

and let $\hat{u}_1, \hat{u}_2$ be an orthonormal basis for the encounter plane normal to $\hat{v}_{\text{rel}}$. Stacking these basis vectors as rows of

$$T = \begin{bmatrix} \hat{u}_1^\top \\ \hat{u}_2^\top \end{bmatrix} \in \mathbb{R}^{2 \times 3} \quad (29)$$

gives the projected relative-position mean and covariance

$$\mu_{Y,k} = T\mu_{\text{rel}}, \quad \Sigma_{Y,k} = T\Sigma_{\text{rel}}T^\top \quad (30)$$

and the collision probability associated with belief b_k is

$$P_c(b_k) = \int_{\|y\| \leq \rho} \mathcal{N}(y; \mu_{Y,k}, \Sigma_{Y,k}) dy \quad (31)$$

where ρ is the sum of the primary and secondary hard-body radii [23, 24].

We evaluate this integral using the fixed-order Gauss–Chebyshev quadrature formulation described by Elrod [31] implemented in NASA CARA’s software toolbox [27]. In contrast to Chan’s series formulation, which obtains the anisotropic case through an approximation to an isotropic formulation, the quadrature directly evaluates the one-dimensional integral associated with the general two-dimensional Gaussian encounter-plane distribution [24]. This is well suited to the eccentric, potentially rotated encounter-plane covariance ellipses produced by our sensor model in Section 4.2.