

Optimized and Trusted Collision Avoidance for Unmanned Aerial Vehicles using Approximate Dynamic Programming

Zachary Sunberg, Mykel Kochenderfer, and Marco Pavone

Abstract—Safely integrating unmanned aerial vehicles into civil airspace is contingent upon development of a trustworthy collision avoidance system. This paper proposes an approach whereby a parameterized resolution logic that is considered *trusted* for a given range of its parameters is adaptively tuned online. Specifically, to address the potential conservatism of the resolution logic with static parameters, we present a dynamic programming approach for adapting the parameters dynamically based on the encounter state. We compute the adaptation policy offline using simulation-based approximate dynamic programming which can handle the high dimensionality of the problem. Numerical experiments show that this approach improves safety and operational performance compared to the baseline resolution logic, while retaining trustworthiness.

I. INTRODUCTION

As unmanned aerial vehicles (UAVs) move toward full autonomy, it is vital that they be capable of effectively responding to anomalous events, such as the intrusion of another aircraft into the vehicle’s flight path. Minimizing collision risk for aircraft in general, and UAVs in particular, is challenging for a number of reasons. First, avoiding collision requires planning in a way that accounts for the large degree of uncertainty in the future paths of the aircraft. Second, the planning process must balance the competing goals of ensuring safety and avoiding disruption of normal operations. Many approaches have been proposed to address these challenges [1]–[13].

At present, there are two fundamentally different approaches to designing a conflict resolution system. The first approach focuses on inspiring confidence and trust in the system by making it as simple as possible for government regulators and vehicle operators to understand. This is accomplished by constructing the algorithm using hand-specified rules. Several algorithms that fit this paradigm have been proposed, and research toward formally verifying their safety-critical properties is underway [3]–[6]. In this paper, we will refer to such algorithms as *trusted resolution logics* (TRLs). Many TRLs have a number of parameters that determine how conservatively the system behaves, a feature that will be exploited later in this paper.

The second approach focuses on optimizing performance. This entails the offline or online computation of a “best” response action. Dynamic programming is widely used for this task [7]–[12]. Conflict resolution systems designed using this approach will be referred to as *directly optimized* systems. Unfortunately, even if a conflict resolution system performs

well in simulation, government regulators and vehicle operators will often (and sometimes rightly) be wary of trusting its safety due to perceived complexity and unpredictability. Even in the best case, such a system would require expensive and time-consuming development of tools for validation as was the case for the recently developed replacement for the traffic alert and collision avoidance system (TCAS) [12].

In this paper, we propose a conflict resolution strategy that combines the strengths of these two approaches. The key idea is to use dynamic programming to find a policy that actively adjusts the parameters of a TRL to improve performance. If this TRL is trusted and certified for a range of parameter values, then the optimized version of the TRL should also be trusted and more easily certifiable.

We test the new approach in a scenario containing a UAV equipped with a perfect (noiseless) sensor to detect the state of an intruder and a simple TRL to resolve conflicts. This TRL, illustrated in Figure 1, determines a path that does not pass within a specified separation distance, D , of the intruder given that the intruder maintains its current heading (except in cases where no such path exists or when the TRL’s heading resolution is too coarse to find such a path). To account for uncertainty in the intruder’s flight path, if D takes a constant value, say \bar{D} , the value must be very large to ensure safety, but this may cause unnecessary departures from normal operation. To overcome this limitation, our approach is to compute an optimized policy $\tilde{\pi}$ that specifies a time varying separation distance, D_t , based on the encounter state, with the goal of ensuring safety without being too conservative.

The problem of dynamically selecting D_t can be formulated as a Markov decision process (MDP). Online solution of $\tilde{\pi}$ using an algorithm such as Monte Carlo Tree Search (MCTS) [14] would be conceptually straightforward, but would require significant computing power onboard the vehicle, and it would be difficult and time-consuming to rigorously certify the implementation of MCTS due to its reliance on pseudo-random number generation. The key contribution of this paper is to devise an *offline* approach to compute $\tilde{\pi}$. Offline optimization of the policy is difficult because of the size of the state space of the MDP, which is the Cartesian product of the continuous state spaces of the UAV and the intruder. To overcome this challenge, we devise a linear value function approximation scheme that exploits the structure of the state space, and we optimize the value function using simulation-based approximate dynamic programming. The policy is then encoded using an approximate post-decision state value function. Armed with this post-decision value function, the UAV can easily extract the optimal action for




Fig. 1. Trusted Resolution Logic. The near mid air collision exclusion zone (circle with radius D) moves with the intruder through time, but is shown here only at the time of closest approach. The TRL (Algorithm 1) finds a straight trajectory close to the nominal path that avoids this zone.

the current state online by evaluating each possible action.

The remainder of the paper is organized as follows: Section II describes the MDP model of an encounter between two UAVs, Section III describes a solution approach based on approximate dynamic programming, and Section IV contains numerical results followed by conclusions in Section V.

II. PROBLEM FORMULATION

The problem of avoiding an intruder using an optimized TRL is formulated as an MDP, referred to as the *encounter* MDP. An encounter involves two aerial vehicles flying in proximity. The first is the vehicle for which we are designing the resolution logic, which will be referred to as the “own UAV” (or simply “UAV”). The second is the intruder vehicle, which may be manned or unmanned and will be referred to as the “intruder”. Specifically, the encounter MDP is defined by the tuple (S, A, T, R) , which consists of

- The state space, S : The state of the encounter, s , consists of (1) the state of the own UAV, $s^{(o)}$, (2) the state of the intruder, $s^{(i)}$, and (3) a boolean variable dev . The variable dev is set to true if the UAV has deviated from its nominal course. Collectively, the state is given by the triple

$$s = \left(s^{(o)}, s^{(i)}, \text{dev} \right). \quad (1)$$

The state components $s^{(o)}$ and $s^{(i)}$ are specified in Section II-B. To model termination, the state space S includes a “dummy” termination state, denoted with s_{term} .

- The action space, A : The actions are the possible values for the separation distance $D \in \mathbb{R}_{\geq 0}$ used in the baseline TRL (see Figure 1).
- The state transition probability density function, $T : S \times A \times S \rightarrow \mathbb{R}$: The value $T(s, D, s')$ is the probability density of transitioning to state s' given that the separation parameter D is used within the TRL at state s . This function is implicitly defined by a generative model that consists of a state transition function $F(\cdot)$ and a stochastic process W modeling uncertainty in the intruder’s behavior. Section II-C contains the definition of $F(\cdot)$, and Section II-B contains the description of W .
- The reward, $R : S \times A \rightarrow \mathbb{R}$: The reward function, defined in Section II-D, rewards reaching a goal and penalizes near mid air collisions (NMACs), deviation from the nominal path, and time outside a goal region.

The remainder of this section motivates the assumptions in this model and describes each component of the encounter MDP.

A. Model Assumptions

Two important simplifying assumptions were made for this initial study. First, the UAV and the intruder move only in the horizontal plane and at constant speed. Modeling horizontal maneuvers is necessary because UAVs will likely have to employ horizontal maneuvers in place of or in addition to the vertical maneuvers that current collision avoidance systems for manned aircraft such as TCAS rely on. This is due to both climb performance limitations and potential regulatory constraints such as the 500ft ceiling for small (<55lb) UAVs in proposed Federal Aviation Administration rules [15]. Constraining altitude and speed simplifies exposition and reduces the size of the state space, keeping the computational burden modest compared with a formulation that includes variable altitude and speed. Extensions to higher fidelity models (e.g., [16]) are possible and left for future research. A higher fidelity model would present a computational challenge, but perhaps not an insurmountable one. For example, the TRL could be extended to handle variable speed and altitude, but the policy that governs the TRL parameters could be optimized only on the most important dimensions of the model (e.g., the horizontal plane).

Second, the intruder dynamics are *independent* of the UAV’s state; in other words, the intruder does not react to the flight path of the UAV. The cooperative setting where both the UAV and the intruder are equipped with a collision avoidance system (CAS) [8], [13], [17] is left for future research.

B. Vehicle states and dynamics

This paper uses a very simple discrete-time model of an encounter between two aerial vehicles. Both the UAV’s and the intruder’s states consist of the horizontal position (x, y) and heading ψ , that is

$$s^{(o)} = \left(x^{(o)}, y^{(o)}, \psi^{(o)} \right), \quad s^{(i)} = \left(x^{(i)}, y^{(i)}, \psi^{(i)} \right), \quad (2)$$

where, as usual, the superscripts (o) and (i) refer to the UAV and intruder quantities, respectively.

Specifically, the intruder makes small random turns, and its flight path is governed by the equations

$$x_{t+1}^{(i)} = \begin{cases} x_t^{(i)} + v^{(i)} \cos(\psi_t^{(i)}) \Delta t & \text{if } w_t = 0 \\ x_t^{(i)} + v^{(i)} \frac{\sin(\psi_t^{(i)} + w_t \Delta t) - \sin(\psi_t^{(i)})}{w_t} & \text{otherwise} \end{cases}$$

$$y_{t+1}^{(i)} = \begin{cases} y_t^{(i)} + v^{(i)} \sin(\psi_t^{(i)}) \Delta t & \text{if } w_t = 0 \\ y_t^{(i)} - v^{(i)} \frac{\cos(\psi_t^{(i)} + w_t \Delta t) - \cos(\psi_t^{(i)})}{w_t} & \text{otherwise} \end{cases}$$

$$\psi_{t+1}^{(i)} = \psi_t^{(i)} + w_t,$$

where Δt is the time step, $v^{(i)}$ is the constant intruder velocity, and w_t is a stochastic disturbance that captures uncertainty in the turning rate. We let W denote the stochastic process $\{w_t : t \in \mathbb{N}\}$. The random variables w_t in W are

assumed independent and identically normally distributed with zero mean and a specified standard deviation, σ_ψ . Subsequently, the intruder dynamics will be collectively referred to as $f^{(i)}(s_t^{(i)}, w_t)$.

The dynamics of the UAV are simplified conventional fixed wing aircraft dynamics, with a single input, namely the roll angle $\phi^{(o)}$. We assume that the roll dynamics are fast compared to the other system dynamics, so that the roll angle $\phi^{(o)}$ may be directly and instantaneously commanded by the control system. This assumption avoids the inclusion of roll dynamics, which would increase the size of the state space. The dynamical equations for the UAV are

$$\begin{aligned} \dot{\psi}_t^{(o)} &= \frac{g \tan \phi_t^{(o)}}{v^{(o)}} \\ x_{t+1}^{(o)} &= \begin{cases} x_t^{(o)} + v^{(o)} \cos(\psi_t^{(o)}) \Delta t & \text{if } \dot{\psi}_t^{(o)} = 0 \\ x_t^{(o)} + v^{(o)} \frac{\sin(\psi_t^{(o)} + \dot{\psi}_t^{(o)} \Delta t) - \sin(\psi_t^{(o)})}{\dot{\psi}_t^{(o)}} & \text{otherwise} \end{cases} \\ y_{t+1}^{(o)} &= \begin{cases} y_t^{(o)} + v^{(o)} \sin(\psi_t^{(o)}) \Delta t & \text{if } \dot{\psi}_t^{(o)} = 0 \\ y_t^{(o)} - v^{(o)} \frac{\cos(\psi_t^{(o)} + \dot{\psi}_t^{(o)} \Delta t) - \cos(\psi_t^{(o)})}{\dot{\psi}_t^{(o)}} & \text{otherwise} \end{cases} \\ \psi_{t+1}^{(o)} &= \psi_t^{(o)} + \dot{\psi}_t^{(o)} \Delta t, \end{aligned}$$

where $v^{(o)}$ is the constant velocity of the UAV. The performance of the UAV is limited by a maximum bank angle

$$|\phi_t^{(o)}| \leq \phi_{\max}. \quad (3)$$

The UAV dynamics will be collectively referred to as $f^{(o)}(s_t^{(o)}, \phi_t^{(o)})$.

C. UAV control system and transition function

The control system for the UAV consists of the TRL and a simple controller that commands a bank angle to track the course determined by the TRL. The TRL determines a heading angle, $\psi_{\text{resolution}}^{(o)}$, that is close to the heading to the goal, $\psi_{\text{goal}}^{(o)}$, and that will avoid future conflicts with the intruder given that the intruder maintains its current heading. This is shown in Figure 1. In order to choose a suitable heading, many candidate headings, denoted as $\psi_{\text{cand}}^{(o)}$, are evaluated.

Given an initial state s , a candidate heading, $\psi_{\text{cand}}^{(o)}$, for the UAV, and that both vehicles maintain their heading, the distance between the vehicles at the time of closest approach is a simple analytical function. Specifically, consider the distance $d(s, \psi_{\text{cand}}^{(o)}, \tau)$ between the vehicles τ time units in the future, i.e.,

$$d(s, \psi_{\text{cand}}^{(o)}, \tau) = \sqrt{\Delta x(\tau)^2 + \Delta y(\tau)^2}, \quad (4)$$

where

$$\begin{aligned} \Delta x(\tau) &= x^{(i)} - x^{(o)} + \tau v^{(i)} \cos(\psi^{(i)}) - \tau v^{(o)} \cos(\psi_{\text{cand}}^{(o)}), \\ \Delta y(\tau) &= y^{(i)} - y^{(o)} + \tau v^{(i)} \sin(\psi^{(i)}) - \tau v^{(o)} \sin(\psi_{\text{cand}}^{(o)}). \end{aligned}$$

The minimum distance between the two vehicles is analytically found by setting the time derivative of $d(s, \psi_{\text{cand}}^{(o)}, \tau)$

Algorithm 1 Trusted Resolution Logic

Input: Encounter state s , desired separation distance D

Output: Resolution heading angle $\psi_{\text{resolution}}^{(o)}$

function TRL(s, D)

$\Psi \leftarrow \{\psi^{(o)} + n\pi/N : n \in \{-N, \dots, N\}, N > 0.\}$
 \triangleright range of values for heading

$D^* = \max_{\psi_{\text{cand}}^{(o)} \in \Psi} d_{\min}(s, \psi_{\text{cand}}^{(o)})$

if $D^* < D$ **then** \triangleright conflict inescapable

$\Psi \leftarrow \{\psi_{\text{cand}}^{(o)} \in \Psi : d_{\min}(s, \psi_{\text{cand}}^{(o)}) = D^*\}$

return $\operatorname{argmin}_{\psi_{\text{cand}}^{(o)} \in \Psi} |\psi_{\text{cand}}^{(o)} - \psi_{\text{goal}}^{(o)}|$

else

$\Psi \leftarrow \{\psi_{\text{cand}}^{(o)} \in \Psi : d_{\min}(s, \psi_{\text{cand}}^{(o)}) \geq D\}$

return $\operatorname{argmin}_{\psi_{\text{cand}}^{(o)} \in \Psi} |\psi_{\text{cand}}^{(o)} - \psi_{\text{goal}}^{(o)}|$

to zero. Specifically, the time at which the vehicles are closest is given by

$$\tau_{\min}(s, \psi_{\text{cand}}^{(o)}) = \max\left\{\frac{a+b}{c-2d}, 0\right\}, \quad (5)$$

where

$$\begin{aligned} a &:= -v^{(i)} x^{(i)} \cos(\psi^{(i)}) - v^{(i)} y^{(i)} \sin(\psi^{(i)}) \\ b &:= v^{(o)} x^{(i)} \cos(\psi_{\text{cand}}^{(o)}) + v^{(o)} y^{(i)} \sin(\psi_{\text{cand}}^{(o)}) \\ c &:= v^{(o)2} + v^{(i)2} \cos^2(\psi^{(i)}) + v^{(i)2} \sin^2(\psi^{(i)}) \\ d &:= v^{(o)} v^{(i)} (\cos(\psi^{(i)}) \cos(\psi_{\text{cand}}^{(o)}) + \sin(\psi^{(i)}) \sin(\psi_{\text{cand}}^{(o)})). \end{aligned}$$

The minimum separation distance over all future time is then

$$d_{\min}(s, \psi_{\text{cand}}^{(o)}) := d\left(s, \psi_{\text{cand}}^{(o)}, \tau_{\min}(s, \psi_{\text{cand}}^{(o)})\right). \quad (6)$$

The TRL begins with a discrete set of potential heading values for the UAV. It then determines, for a desired separation distance D , which of those will not result in a collision given that the UAV and intruder maintain their headings. Finally, it selects the value from that set which is closest to $\psi_{\text{goal}}^{(o)}$. The TRL is outlined in Algorithm 1.

Once the TRL has returned the desired heading (denoted as $\psi_{\text{resolution}}^{(o)}$), a low-level controller determines the control input to the vehicle. We write this as

$$\phi_t^{(o)} = c\left(s_t^{(o)}, \psi_{\text{resolution}}^{(o)}\right), \quad (7)$$

where $c(\cdot)$ represents the low level controller.

The closed-loop dynamics for the state variables of the vehicles are then given by

$$s_{t+1}^{(i)} = f^{(i)}\left(s_t^{(i)}, w_t\right) \quad (8)$$

$$s_{t+1}^{(o)} = f^{(o)}\left(s_t^{(o)}, c\left(s_t^{(o)}, \text{TRL}(s_t, D_t)\right)\right). \quad (9)$$

The UAV and the intruder dynamics are coupled *only* through the TRL.

The goal region that the UAV is trying to reach is denoted by S_{goal} . This is the set of all states in S for which $\|(x^{(o)}, y^{(o)}) - (x_{\text{goal}}^{(o)}, y_{\text{goal}}^{(o)})\| \leq D_{\text{goal}}$, where $D_{\text{goal}} > 0$

is a specified goal region radius, and $(x_{\text{goal}}^{(o)}, y_{\text{goal}}^{(o)})$ is the goal center location. A near mid air collision (NMAC) occurs at time t if the UAV and intruder are within a minimum separation distance, $D_{\text{NMAC}} > 0$, that is if $\| (x_t^{(o)}, y_t^{(o)}) - (x_t^{(i)}, y_t^{(i)}) \| \leq D_{\text{NMAC}}$. If the UAV reaches the goal region at some time t , i.e., $s_t \in S_{\text{goal}}$, or if an NMAC occurs, the overall encounter state s transitions to the terminal state s_{term} and remains there. If the UAV performs a turn, `dev` is set to true because the vehicle has now deviated from the nominal straight path to the goal.

Let the state transition function, defined by (8), (9), and the special cases above, be denoted concisely as F so that

$$s_{t+1} = F(s_t, D_t, w_t), \quad (10)$$

where, as stated above, D_t is the action input to the system.

D. Reward

In this paper, we optimize two competing metrics. The first goal is to minimize the risk of a NMAC. As in previous studies (e.g., [7], [9]), this aspect of performance is quantified using the *risk ratio*, the number of near-collisions with the control system divided by the number of near-collisions without the control system. The second goal is to minimize the probability of deviation from the nominal path. This metric was chosen (as opposed to a metric that penalizes the size of the deviation) because in a commercial setting, any deviation from the normal operating plan might have a large cost in the form of disrupting schedules, preventing a mission from being completed, or requiring manual human monitoring. The MDP reward function is designed to encourage the policy towards a solution that performs well with respect to these goals.

Specifically, the utility associated with an encounter is the sum of the stage-wise rewards throughout the entire encounter

$$\sum_{t=0}^{\infty} R(s_t, D_t). \quad (11)$$

In order to meet both goals, the stage-wise reward function is defined as

$$\begin{aligned} R(s_t, D_t) := & -c_{\text{step}} + r_{\text{goal}} \times \text{in_goal}(s_t^{(o)}) \\ & - c_{\text{dev}} \times \text{causes_deviation}(s_t, D_t) \\ & - \lambda \times \text{is_NMAC}(s_t), \end{aligned} \quad (12)$$

for positive constants c_{step} , r_{goal} , c_{dev} , and λ . The first term is a small constant cost accumulated in each step to push the policy to quickly reach the goal. The function `in_goal` indicates that the UAV is within the goal region, so the second term is a reward for reaching the goal. The third term is a penalty for deviating from the nominal path. The function `causes_deviation` returns 1 if the action will cause a deviation from the nominal course and 0 otherwise. It will only return 1 if the vehicle has not previously deviated and `dev` is false, so the penalty may only occur once during an encounter. This behavior makes the inclusion of `dev` in the state necessary. Constants c_{step} , r_{goal} , and c_{dev} represent relative weightings for the terms that incentivize a policy

that reaches the goal quickly and minimizes the probability of deviation. Example values for these constants are given in Section IV. The fourth term is the cost for a collision. The weight λ balances the two performance goals. We heuristically expect there to be a value of λ for which the solution to the MDP meets the desired risk ratio if it is attainable. Bisection, or even a simple sweep of values can be used to find a suitable value. This method of adjusting weights in the reward function has been used previously to analyze the performance of aircraft collision avoidance systems [8], [11].

E. Problem statement

The problem we consider is to find a feedback control policy $\pi^* : S \rightarrow A$, mapping an encounter state s_t into a separation distance D_t , that maximizes the expected reward (11) subject to the system dynamics (10):

$$\begin{aligned} & \underset{\pi}{\text{maximize}} && E_W \left[\sum_{t=0}^{\infty} R(s_t, \pi(s_t)) \right] \\ & \text{subject to} && s_{t+1} = F(s_t, \pi(s_t), w_t), \end{aligned} \quad (13)$$

for all initial states $s_0 \in S$. To make the solution of problem (13) practical, we present an approximate dynamic programming approach that yields a suboptimal policy, $\tilde{\pi}$, in the next section.

III. SOLUTION APPROACH

The solution approach we applied to the problem (13) is an approximate dynamic programming algorithm called approximate value iteration [18]. The value function, V , represents the expected value of the future reward given that the encounter is in state s and an optimal policy will be executed in the future. We approximate V with a linear architecture of the form [18]

$$\tilde{V}(s) = \beta(s)^\top \theta, \quad (14)$$

where the feature function β returns a vector of N_β feature values, and $\theta \in \mathbb{R}^{N_\beta}$ is a vector of weights [18]. At each step of value iteration, the weight vector θ is fitted to the results of a large number of single-step simulations by solving a linear least-squares problem. After value iteration has converged, \tilde{V} is used to estimate a *post-decision state* value function, \tilde{V}_q , which is also approximated using a linear combination of features. The policy is extracted online in real time by selecting the action that results in the post-decision state that has the highest value according to \tilde{V}_q . The choice of working with post-decision states will be discussed in Section III-B.

A. Approximate Value Iteration

The bulk of the computation is carried out offline before vehicle deployment using simulation. Specifically, the first step is to estimate the optimal value function for problem (13) using value iteration. Value iteration is the successive application of the Bellman operator, \mathcal{B} , to the value function

$$V_{k+1}(s) = \mathcal{B}[V_k](s) = \max_{a \in A} \left\{ R(s, a) + E_{s'} [V_k(s') | s, a] \right\}. \quad (15)$$

The optimal value is the fixed point of this equation [18]. On a continuous state space such as the encounter state space used in this work, the Bellman operator cannot be applied for each of the uncountably infinite number of states, so an approximation must be used. In this paper we adopt projected value iteration [18], which uses a finite number of parameters to approximate the value function. Each successive approximation, \tilde{V}_k , is the result of the Bellman operation projected onto a linear subspace with respect to the Euclidean norm, that is

$$\tilde{V}_{k+1}(s) = \Pi B[\tilde{V}_k](s), \quad (16)$$

where the operator Π is a projection onto the linear subspace

$$\Phi = \{B\theta : \theta \in \mathbb{R}^{N_\beta}\}, \quad (17)$$

where B is a matrix of which each row is the feature vector of a state (the number of rows is uncountably infinite) [18]. Intuitively, the value function V is approximated within the subspace Φ spanned by N_β basis functions, the columns of Φ (see equation (14)).

To perform the approximate value iteration (16), we resort to Monte Carlo simulations. Specifically, for each iteration, N_{state} states are uniformly randomly selected. If the states lie within the grids used in the feature function (see Section III-D) the sample is ‘‘snapped’’ to the nearest grid point to prevent approximation errors due to the Gibbs phenomenon [19]. At each sampled state $s^{[n]}$, $n = 1, \dots, N_{\text{state}}$, the stage-wise reward and the expected value of the value function are evaluated for each action a within a discrete approximation of the action space A , denoted by \tilde{A} . The expectation in the right-hand side of (15) is approximated using N_{EV} single-step intruder simulations, each with a randomly generated noise value, w_m , $m = 1, \dots, N_{\text{EV}}$. However, since the own UAV dynamics are deterministic, only one own UAV simulation is needed. The maximum over the set of actions \tilde{A} is selected and stored as the n th entry of a vector v_{k+1} :

$$v_{k+1}[n] := \max_{D \in \tilde{A}} \left\{ R(s^{[n]}, D) + \frac{1}{N_{\text{EV}}} \sum_{m=1}^{N_{\text{EV}}} \beta(F(s^{[n]}, D, w_m))^\top \theta_k \right\},$$

for $n = 1, \dots, N_{\text{state}}$. Here v_{k+1} provides an approximation to the (unprojected) value function V_{k+1} . To project v_{k+1} onto Φ , we compute the weight vector θ_{k+1} by solving the least square optimization problem

$$\theta_{k+1} = \underset{\theta \in \mathbb{R}^{N_\beta}}{\operatorname{argmin}} \sum_{n=1}^{N_{\text{state}}} \left(\beta(s^{[n]})^\top \theta - v_k[n] \right)^2. \quad (18)$$

Let \tilde{B} be a matrix so that each row corresponds to the feature vector of $s^{[n]}$, that is $\tilde{B}[n, :] = \beta(s^{[n]})$ (the size of \tilde{B} is $N_{\text{state}} \times N_\beta$). The optimization problem above can then be compactly written as

$$\theta_{k+1} = \underset{\theta \in \mathbb{R}^{N_\beta}}{\operatorname{argmin}} \|\tilde{B}\theta - v_k\|_2^2, \quad (19)$$

which is readily solved by taking the Moore-Penrose pseudo-inverse

$$\theta_{k+1} = \left(\tilde{B}^\top \tilde{B} \right)^{-1} \tilde{B}^\top v_k. \quad (20)$$

Iteration is terminated after a fixed number of steps, N_{VI} , and the resulting weight vector, denoted by θ , is stored for the next processing step (Section III-B).

B. Post Decision Value Function Extraction

For reasons discussed in Section III-C, our second step is to approximate a value function defined over post-decision states [18]. A post-decision state, denoted by q , is defined as a state in S consisting of the own UAV state and dev at one time step into the future and the intruder state at the current time, that is

$$q_t = \left(s_{t+1}^{(o)}, s_t^{(i)}, \text{dev}_{t+1} \right). \quad (21)$$

Correspondingly, let $g : S \times A \rightarrow S$ be the function that maps the current state and action to the post-decision state. In other words, function $g(s_t, D_t)$ returns q_t consisting of

$$\begin{aligned} s_{t+1}^{(o)} &= f^{(o)} \left(s_t^{(o)}, c \left(s_t^{(o)}, TRL(s_t, D_t) \right) \right) \\ s_t^{(i)} &= s_t^{(i)} \\ \text{dev}_{t+1} &= \max\{\text{dev}, \text{causes_deviation}(s_t, D_t)\}. \end{aligned} \quad (22)$$

The approximate value function over post-decision states, \tilde{V}_q , is computed as follows. Let $h : S \times \mathbb{R} \rightarrow S$ be a function that returns the next encounter state given the post decision state and intruder heading noise value, that is $h(q_t, w)$ returns s_{t+1} consisting of

$$s_{t+1}^{(o)} = s_{t+1}^{(o)} \quad (23)$$

$$s_{t+1}^{(i)} = f^{(i)}(s_t^{(i)}, w) \quad (24)$$

$$\text{dev}_{t+1} = \text{dev}_{t+1}, \quad (25)$$

where $(s_{t+1}^{(o)}, s_t^{(i)}, \text{dev}_{t+1})$ are the members of q_t .

The value function \tilde{V}_q can then be expressed in terms of \tilde{V} as

$$\tilde{V}_q(q) = E_w \left[\tilde{V}(h(q, w)) \right], \quad (26)$$

where w denotes, as usual, a random variable with Gaussian normal density. Equation (26) implies that \tilde{V}_q can also be approximated according to a linear architecture

$$\tilde{V}_q(q) = \beta(q)^\top \theta^q, \quad (27)$$

where $\beta(q)$ is the feature vector for post-decision states $q \in S$, and $\theta^q \in \mathbb{R}^{N_\beta}$ is the corresponding weight vector.

Specifically, N_q post decision states are randomly selected using the same method as described in Section III-A and are denoted as $q^{[n]}$, $n = 1, \dots, N_q$. For each sampled state $q^{[n]}$, the expectation in (26) is approximated using N_{EV} single-step simulations. The results are used to solve a least squares optimization problem

$$\theta^q = \underset{\theta \in \mathbb{R}^{N_\beta}}{\operatorname{argmin}} \sum_{n=1}^{N_q} \left(\beta(q^{[n]})^\top \theta - v^q[n] \right)^2, \quad (28)$$

where

$$v^q[n] := \frac{1}{N_{\text{EV}}} \sum_{m=1}^{N_{\text{EV}}} \beta(h(q^{[n]}, w_m))^\top \theta, \quad (29)$$

where w_m , $m = 1, \dots, N_{\text{EV}}$, is a noise value sampled from a random variable in W .

C. Online Policy Evaluation

The first two steps (explained, respectively, in Sections III-A and III-B) are performed offline. The last step, namely policy evaluation, is performed online. Specifically a suboptimal control at any state s is computed as

$$\tilde{\pi}(s) = \operatorname{argmax}_{D \in \tilde{A}} \tilde{V}_q(g(s, D)). \quad (30)$$

Since g (defined in (22)) is a *deterministic* function of a state-action pair, this calculation does not contain any computationally costly or difficult-to-certify operations such as estimating an expectation.

Two comments regarding the motivation for using the post-decision value function are in order. First, if the value functions could be exactly calculated, the post-decision state approach would be equivalent to the more common Q -factor-based approach, wherein a control is computed by solving

$$\pi(s) = \operatorname{argmax}_{D \in \tilde{A}} Q(s, D), \quad (31)$$

and $Q(s, a)$ (the Q -factor) represents the expected value of taking action a in state s and then following an optimal policy [18]. In fact, one can readily show $Q(s, D) = V_q(g(s, D))$. However, when approximations are used, post-decision states provide a more robust (i.e., less susceptible to approximation error) way of deriving control actions [18]. One reason for this is that the cost function is approximated in the space of post-decision states, rather than in the larger space of state-control pairs, and hence the post-decision method is less susceptible to complications with inadequate exploration [18]. Second, the post decision value function is not used for the value iteration portion of the offline solution as it would require full simulations of both the own UAV and the intruder dynamics and, therefore, would be more computationally demanding.

D. Selection of Features

The primary value function approximation features are the interpolation weights for points in a grid [20]. Since a grid defined over the entire six dimensional encounter state space with a reasonable resolution would require far too many points to be computationally feasible, the grid must be focused on important parts of the state space.

Our strategy is to separate features into two groups (along with a constant), specifically

$$\beta(s) = [\beta_{\text{intruder}}(s^{(o)} - s^{(i)}), \beta_{\text{goal}}(s^{(o)}), 1]. \quad (32)$$

The first group, β_{intruder} , captures the features corresponding to a near midair collision and is a function of only the position and orientation of the own vehicle relative to the intruder. The second group, β_{goal} , captures the value of being near the goal and is a function of only the own vehicle state.

Since the domain of β_{intruder} is only three dimensional and the domain of β_{goal} is only two dimensional, relatively fine interpolation grids can be used for value function approximation without requiring a prohibitively large number of features. The β_{intruder} feature group consists of a NMAC indicator function and interpolation weights for a grid (Figure 2) with nodes at regularly spaced points along

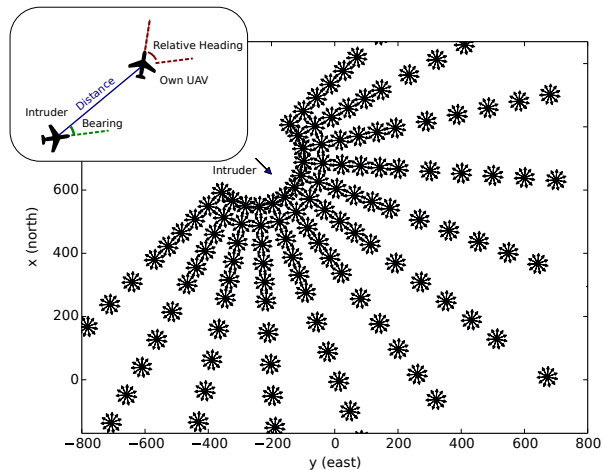


Fig. 2. Intruder interpolation grid for β_{intruder} , visualized when the intruder is located at (700 m, -250 m) at heading 135. The top left inset shows the variables used. In the main plot, each of the small arrows represents a grid point. At each of the point locations, there are twelve small arrows radiating out. Each arrow represents a different own UAV heading.




Fig. 3. Two-dimensional slice of the approximate optimal value function. Each pixel on this image represents the value function evaluated with the intruder at that position pointed directly north. The intruder is located at (700 m, -250 m) pointed at heading 135 as indicated by the arrow. The goal is at (1000 m, 0)

the following three variables: (1) the distance between the UAV and intruder, (2) the bearing from the intruder to the UAV, and (3) the relative heading between the vehicles. The β_{goal} vector consists of a goal indicator function, the distance between the UAV and the goal, and interpolation weights for a grid with nodes regularly spaced along the distance between the UAV and the goal and the absolute value of the bearing to the goal from the UAV. The total number of features is $N_{\beta} = 1813$.

Figure 3 shows a single two-dimensional “slice” of the full six-dimensional optimized value function. As expected, there is a low value region in front of and to the south of the intruder and an increase in the value near the goal.

Figure 4 shows a slice of an optimized policy. When multiple actions result in the same post decision state value, the least conservative action is chosen, so the policy yields the lowest value of D (500ft \approx 152.4m) on most of the




Fig. 4. Two-dimensional slice of the optimized policy for TRL parameter D . Each pixel in the image is the policy evaluated with the own UAV at that location pointed directly north. The intruder is located at (700 m, -250 m) and pointed at heading 135 as indicated by the arrow. Each color represents a different value for D indicated by the number in the key on the right.

state space. Because the own UAV is pointed north, the policy is conservative in a region in front of and to the south of the intruder. The band corresponding to small D that stretches across the middle of the conservative region (from (700 m, -250 m) to (-100 m, 200 m)) is present because all values of D result in the same post decision value.

IV. RESULTS

This section presents results from numerical experiments that illustrate the effectiveness of our approach. The experiments are designed to compare the three approaches discussed in Section I: the “static TRL” approach, the “directly optimized” approach, and the new “optimized TRL” approach proposed in this paper. Each of these approaches is represented by a different control law that specifies a bank angle control to the own UAV. Specifically, the static TRL law uses the TRL described in Algorithm 1 with a constant value for the separation distance (denoted with \bar{D}). To represent the directly optimized approach, an optimized value function approximation and policy are generated using the method described in Section III, except that the policy *directly determines the bank angle*, $\phi_t^{(o)}$, *instead of specifying D* . The action space for this policy is $\{-\phi_{\max}, -\phi_{\max}/2, 0, \phi_{\max}/2, \phi_{\max}\}$. The third optimized TRL law works as described in Sections II and III, dynamically assigning a value of D to the TRL based on the state and using $c(\cdot)$ to assign the bank angle. The action space for the underlying D policy is $\hat{A} = \{D_{\text{NMAC}}, 1.5D_{\text{NMAC}}, 2D_{\text{NMAC}}, 3D_{\text{NMAC}}, 4D_{\text{NMAC}}\}$.

The various parameters used in the numerical trials are listed in Table I. The control policies are evaluated by executing them in a large number of complete (from $t = 0$ to the end state) encounter simulations. The same random numbers used to generate intruder noise were reused across all collision avoidance strategies to ensure fairness of comparisons. In each of the simulations, the own UAV starts pointed north at position (0, 0) in a north-east coordinate system with the goal at (1000 m, 0).

TABLE I
PARAMETERS FOR NUMERICAL EXPERIMENTS

| Description | Symbol | Value |
|---|--------------------|--------|
| Own UAV speed | $v^{(o)}$ | 30 m/s |
| Maximum own UAV bank angle | ϕ_{\max} | 45° |
| Intruder speed | $v^{(i)}$ | 60 m/s |
| Intruder turn rate standard deviation | σ_{ψ} | 10°/s |
| Near mid air collision radius | D_{NMAC} | 500 ft |
| Step cost | c_{step} | 1 |
| Reward for reaching goal | r_{goal} | 100 |
| Cost for deviation | c_{dev} | 100 |
| Step simulations for expectation estimate | N_{EV} | 20 |
| Single step simulations per round of value iteration (optimized TRL) | N_{state} | 10,000 |
| Single step simulations per round of value iteration (directly optimized) | N_{state} | 50,000 |
| Number of value iteration rounds | N_{VI} | 35 |
| Single step simulations for post decision value function extraction | N_q | 50,000 |

The conservativeness of each control law is characterized by counting the number of deviations from the nominal path in 10,000 simulations. The intruder initial conditions for these simulations are randomly generated and lie between 800 m and 1500 m from the center point of the encounter area at (500 m, 500 m). The intruder’s initial heading is within 135° of the direction from its initial position to the center point. Of these simulations, 1009 result in a NMAC if the UAV follows its nominal path, but for most a deviation is not necessary to avoid the intruder.

The risk ratio is estimated using a separate set of 10,000 simulations. Each of these simulations has an initial condition in the same region described above, but initial conditions and noise trajectories are chosen by filtering random trials so that each of the simulations *will result in a NMAC if the own UAV follows its nominal path*. The risk ratio estimate for a policy is simply the fraction of these simulations that result in a NMAC when the policy is executed on them.

Figure 5 shows the Pareto optimal frontiers for the different control laws. Each curve is generated by using various values of λ in the reward function of the MDP, or by using various static values of \bar{D} in the static TRL case. It is clear that the optimization improves the performance of the TRL. For example, if the desired risk ratio is 5%, interpolation between data points suggests that the optimized TRL policy will cause approximately 10% fewer deviations than the static TRL policy.

The directly optimized approach performs better than both TRL policies. This is not unexpected since the directly optimized policy is not limited to trajectories that the TRL deems to be safe. However, precisely because the directly optimized policy is not limited to trajectories guaranteed safe by hand-specified rules, it is much more difficult to convince people to trust it. These results thus allow us to estimate the performance *price* of using a trusted resolution logic rather than an untrusted one. Since the optimized TRL causes approximately 18% more deviations than the directly optimized policy at a risk ratio of 0.05, we may say the price of the trustworthy properties of the TRL at a desired risk ratio of 0.05 is an 18% increase in deviations. The new optimized TRL approach has reduced this from a price of a

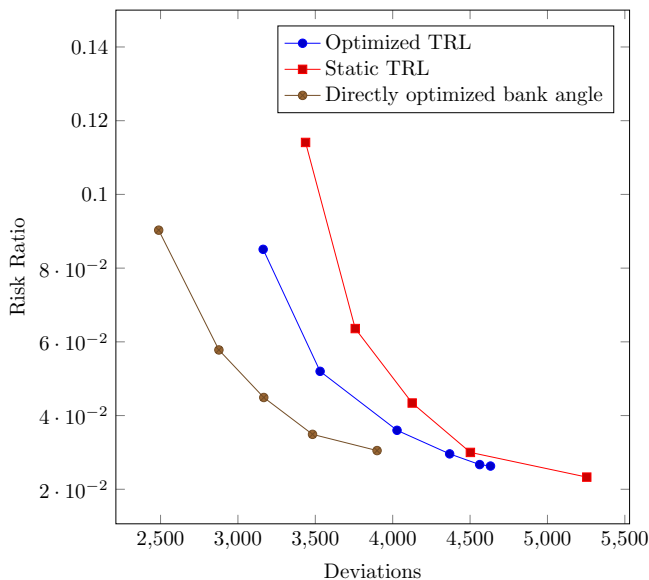


Fig. 5. Control policy comparison. The horizontal axis variable is the number of deviations in 10,000 encounter simulations. The values of λ used to generate the datapoints are 100, 316, 1000, 3160, 10^4 , and 3.16×10^4 for the optimized TRL policy and 300, 500, 700, 1000, and 1500 for the directly optimized approach. The values of \bar{D} for the static TRL policy are 250 m, 300 m, 350 m, 400 m, and 500 m.

31% increase in deviations without optimization.

The software used to simulate these experiments was written in the Julia programming language, and is freely available at <https://github.com/zsunberg/UASEncounter>.

V. CONCLUSION

This paper presented a method for improving the performance of trusted conflict resolution logic for an unmanned aerial vehicle through approximate dynamic programming. A linear value function approximation based on interpolation grids and other features produces policies that improve the performance of the resolution logic without undermining the trust placed in it. Specifically, simulation experiments show that this optimization approach is able to decrease the number of deviations from the nominal path without increasing the risk ratio. Comparison with a directly optimized approach that is more difficult to trust can help to quantify the price of the gain in trustworthiness that comes from using the trusted resolution logic. The optimization approach proposed here is capable of reducing that price.

There is much potential for future work on this problem. The current formulation models uncertainty in the intruder's actions, but sensor uncertainty should also be taken into account. If sensor uncertainty is added to the problem, it becomes a partially observable Markov decision process. One promising solution approach for the partially observable version of this problem is Monte Carlo value iteration [21]. This approach has been applied to collision avoidance for unmanned systems before [11], but it would be difficult to certify as a direct control system. Monte Carlo value iteration used in conjunction with a TRL has potential to handle both uncertainty in intruder behavior and sensor measurement noise while being easily certifiable. Also, as mentioned in Section II-A, there is potential to expand this work to more complicated scenarios with multiple intruders,

intruders equipped with a CAS, higher fidelity models, or different sets of assumptions.

VI. ACKNOWLEDGMENTS

This material is based on work supported in part by the National Science Foundation, Grant No. 1252521 and in part by the Office of Naval Research, Science of Autonomy Program, under Contract N00014-15-1-2673.

REFERENCES

- [1] J. K. Kuchar and L. C. Yang, "A review of conflict detection and resolution modeling methods," *IEEE Transactions on Intelligent Transportation Systems*, vol. 1, no. 4, pp. 179–189, December 2000.
- [2] H. Y. Ong and M. J. Kochenderfer, "Short-term conflict resolution for unmanned aircraft traffic management," in *Digital Avionics Systems Conference*, Sep. 2015.
- [3] R. Bach, C. Farrell, and H. Erzberger, "An algorithm for level-aircraft conflict resolution," NASA, Tech. Rep. CR-2009-214573, 2009.
- [4] G. Hagen, R. Butler, and J. Maddalon, "Stratway: A modular approach to strategic conflict resolution," in *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, Virginia Beach, 2011.
- [5] H. Herencia-Zapana, J.-B. Jeannin, and C. Muñoz, "Formal verification of safety buffers for safe-based conflict detection and resolution," in *International Congress of the Aeronautical Sciences*, 2010.
- [6] A. Narkawicz, C. Muñoz, and G. Dowek, "Provably correct conflict prevention bands algorithms," *Science of Computer Programming*, vol. 77, no. 10–11, pp. 1039–1057, 2012.
- [7] S. Temizer, M. J. Kochenderfer, L. P. Kaelbling, T. Lozano-Perez, and J. K. Kuchar, "Collision avoidance for unmanned aircraft using markov decision processes," in *AIAA Guidance, Navigation, and Control Conference*, 2010.
- [8] M. J. Kochenderfer and J. P. Chryssanthacopoulos, "Robust airborne collision avoidance through dynamic programming," MIT Lincoln Laboratory, Tech. Rep. ATC-371, 2011.
- [9] M. J. Kochenderfer, J. P. Chryssanthacopoulos, L. P. Kaelbling, and T. Lozano-Perez, "Model-based optimization of airborne collision avoidance logic," MIT Lincoln Laboratory, Tech. Rep. ATC-360, 2010.
- [10] L. P. Kaelbling and T. Lozano-Perez, "Finding aircraft collision-avoidance strategies using policy search methods," MIT-CSAIL, Tech. Rep. TR-2009-043, 2009.
- [11] H. Bai, D. Hsu, M. J. Kochenderfer, and W. S. Lee, "Unmanned aircraft collision avoidance using continuous-state POMDPs," in *Robotics: Science and Systems*, vol. 7, 2012.
- [12] J. E. Holland, M. J. Kochenderfer, and W. A. Olson, "Optimizing the next generation collision avoidance system for safe, suitable, and acceptable operational performance," in *USA/Europe Air Traffic Management Research and Development Seminar*, 2013.
- [13] E. J. Rodríguez-Seda, "Decentralized trajectory tracking with collision avoidance control for teams of unmanned vehicles with constant speed," in *American Control Conference*, 2014.
- [14] A. Couëtoux, J.-B. Hoock, N. Sokolovska, O. Teytaud, and N. Bonnard, "Continuous upper confidence trees," in *Learning and Intelligent Optimization*. Springer, 2011, pp. 433–445.
- [15] Federal Aviation Administration, "Small uas notice of proposed rulemaking," February 2015. [Online]. Available: <https://www.faa.gov/regulations.policies/rulemaking/recently-published/media/2120-AJ60-NPRM.2-15-2015-joint-signature.pdf>
- [16] M. J. Kochenderfer, M. W. M. Edwards, L. P. Espindle, J. K. Kuchar, and J. D. Griffith, "Airspace encounter models for estimating collision risk," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 33, no. 2, pp. 487–499, 2010.
- [17] C. Tomlin, G. J. Pappas, and S. S. Sastry, "Conflict resolution for air traffic management: A study in multiagent hybrid systems," *IEEE Transactions on Automatic Control*, vol. 43, no. 4, pp. 509–21, 1998.
- [18] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.
- [19] J. Foster and F. B. Richards, "The Gibbs phenomenon for piecewise-linear approximation," *The American Mathematical Monthly*, vol. 98, no. 1, pp. 47–49, January 1991.
- [20] S. Davies, "Multidimensional triangulation and interpolation for reinforcement learning," in *Conf. on Neural Information Processing Systems*, 1996, pp. 1005–1011.
- [21] H. Bai, D. Hsu, W. S. Lee, and V. A. Ngo, "Monte Carlo value iteration for continuous-state POMDPs," in *Algorithmic Foundations of Robotics IX*. Springer, 2011, pp. 175–191.