

Coupled Path-Attitude Optimization for Low-Observable Flight in Complex Terrain via Differentiable Neural Surrogates

Minjae Kim^{a,b}, Chang-Kyung Ryoo^{a,b,c} and Jong-Han Kim^{a,b,c,*}

^aDepartment of Aerospace Engineering, Inha University, 100 Inha-ro, Michuhol-gu, Incheon 22212, Republic of Korea

^bProgram in Aerospace Systems Convergence, Inha University, 100 Inha-ro, Michuhol-gu, Incheon 22212, Republic of Korea

^cAerospace Systems Research Institute, Inha University, 100 Inha-ro, Michuhol-gu, Incheon 22212, Republic of Korea

ARTICLE INFO

Keywords:

Trajectory optimization
Low observability
Terrain masking
Radar detection
Sequential convex programming
Surrogate model
Weapon engagement zone

ABSTRACT

This paper presents a coupled path-attitude optimization framework for low-observable UAV flight in complex terrain under radar and surface-to-air missile (SAM) threats. The method combines terrain masking, radar signal-to-noise ratio (SNR) reduction, and dynamic weapon engagement zone (WEZ) avoidance in a single trajectory design problem. To represent the strong nonlinear dependence of detection risk on vehicle position, attitude, and terrain geometry, multilayer perceptron surrogates are constructed for the radar SNR and a terrain exposure index. The surrogate models provide differentiable approximations that can be embedded in a sequential convex programming (SCP) loop, while automatic differentiation is used to obtain the Jacobians required for linearization. To reduce the impact of false-negative visibility predictions near exposure boundaries, a conditional value-at-risk (CVaR)-based safety margin is introduced in the terrain exposure model. The resulting optimization problem includes nonlinear flight dynamics, surrogate-based detection metrics, terrain clearance, and dynamic WEZ constraints. Simulation studies in a mountainous environment show that the framework generates feasible trajectories that avoid the prescribed WEZs and exploit terrain occlusion when available. In exposed segments, the optimizer also selects vehicle attitudes associated with lower surrogate-predicted SNR, indicating that position and attitude can be coordinated to reduce the predicted radar detection metric within the proposed surrogate-based formulation. Weight-sweep results further illustrate the trade-off between terrain masking preference and SNR reduction.

1. Introduction

In recent anti-access/area-denial (A2/AD) environments, in which layered surveillance and weapon systems are deployed to restrict aircraft access to a protected region, securing the survivability of aircraft or unmanned aerial vehicles (UAVs) has become a central problem in mission planning. In particular, trajectory design that reduces detectability from integrated air-defense sensors, such as radar and electro-optical/infrared (EO/IR) systems, strongly influences UAV operational effectiveness. In mountainous terrain, nap-of-the-earth (NOE) flight can exploit terrain occlusion to reduce radar line-of-sight (LOS) exposure, where LOS denotes the direct geometric visibility path between the radar and the UAV [1]. Under surface-to-air missile (SAM) threats, however, relying only on post-launch evasive maneuvers can be insufficient because the available reaction time, maneuvering space, and low-altitude clearance margin may be severely limited after a missile launch is detected. Therefore, survivability-oriented planning should be performed before engagement by shaping the trajectory to exploit terrain masking and to avoid weapon engagement zones (WEZs), which represent the effective regions in which a SAM system can engage the vehicle. These considerations motivate trajectory optimization methods that jointly account for terrain masking, radar detectability, and dynamically changing WEZ constraints [2, 3].

Compared with purely reactive evasion, trajectory optimization provides a proactive design mechanism because terrain geometry, sensor visibility, vehicle dynamics, and threat constraints can be considered simultaneously over the entire mission horizon. Various trajectory optimization studies have been conducted to address these complex threats, reflecting a broader shift in the field from classical graph-theoretic algorithms toward advanced AI-driven and optimization-based path planning [4]. An early study by Pelosi et al. [5] proposed a terrain masking path planning method based on a 3D grid search; however, spatial discretization makes it difficult to reflect 6-DoF dynamic constraints and changes in detectability due to the vehicle's attitude. To overcome this, Xu et al. [6] presented a trajectory design method that lowers detection probability by coupling the anisotropy of radar cross section (RCS) with the aircraft's attitude. Nevertheless, because this approach fits highly nonlinear RCS data with polynomials and uses metaheuristic algorithms like particle swarm optimization (PSO), it suffers from heavy computational loads and is difficult to apply to an efficient optimization framework. Recently, studies using reinforcement learning [7] or dynamic cardioid WEZ avoidance [8] have been reported. Concurrently, the integration of deep neural networks into optimal control has emerged as a promising approach to alleviate computational burdens and efficiently handle complex spatial constraints [9, 10, 11]. Moreover, recent advances have demonstrated the efficacy of transforming highly nonlinear kinodynamic UAV constraints into convex optimization problems to enable rapid trajectory generation [12]. However, learning-based models

*Corresponding author: J.-H. Kim

✉ jonghank@inha.ac.kr (J. Kim)

🌐 <https://jonghank.github.io> (J. Kim)

ORCID(s): 0000-0002-9030-0490 (J. Kim)

often lack mathematical safety guarantees, and there remains a technical gap in integrating advanced stealth and terrain masking strategies into a single, rigorous optimization framework.

To address these limitations, this paper proposes a coupled path-attitude design framework that integrates terrain masking and attitude-based low observability within a computationally tractable optimization setting. Unlike path-only planning, which primarily exploits terrain masking and geometric separation from threats, the proposed formulation simultaneously optimizes the UAV path and attitude. This coupling allows the UAV to exploit terrain occlusion when available and to select attitudes associated with lower surrogate-predicted SNR when exposure is unavoidable. Radar signal-to-noise ratio (SNR) from the radar equation and a terrain exposure index based on skyline elevation are used as survivability metrics. Because both quantities vary nonlinearly with the flight state, they are approximated using differentiable multilayer perceptron (MLP) surrogate models and embedded in a sequential convex programming (SCP) loop [13, 14]. Automatic differentiation is used to compute the Jacobians required for linearization. In addition, a conditional value-at-risk (CVaR)-based safety margin [15] is applied to the exposure model to conservatively reduce the effect of false-negative exposure predictions near visibility transition boundaries. The resulting formulation combines nonlinear flight dynamics, surrogate-based survivability metrics, and dynamic WEZ avoidance constraints in a single optimization problem.

The main contributions of this paper are summarized as follows:

- *Coupled survivability formulation:* A UAV trajectory design problem is formulated to jointly account for terrain masking, radar SNR reduction, and dynamic WEZ avoidance through coupled path and attitude decisions.
- *Differentiable surrogate modeling:* MLP-based surrogate models are developed for the terrain exposure index and radar SNR so that these nonlinear survivability metrics can be embedded in an SCP framework with automatic differentiation.
- *Conservative treatment of visibility uncertainty:* A CVaR-based safety margin is introduced to reduce sensitivity to false-negative exposure predictions near visibility transition boundaries.
- *Simulation-based demonstration:* Numerical studies in mountainous terrain illustrate how the proposed framework trades terrain masking against attitude-based SNR reduction while maintaining dynamic WEZ avoidance.

2. Vehicle, Radar, and Threat Modeling

The following basic modeling assumptions define the scope of this study:

1. The UAV is modeled as a pseudo 6-DoF point mass in the NED frame. Translational motion and yaw are optimized, while roll and pitch are reconstructed from the thrust vector and yaw angle.
2. The aerodynamic model includes only quadratic drag. Wind, atmospheric variation, actuator dynamics, and detailed rotor/propulsion dynamics are not explicitly modeled.
3. Radar detectability is evaluated using a single-radar SNR model based on a precomputed attitude-dependent RCS database, and terrain exposure is computed from deterministic LOS geometry over a fixed terrain map.
4. The SAM threat is represented by a horizontal cardioid-shaped WEZ, and WEZ avoidance is enforced as a deterministic state constraint.

These assumptions are introduced to keep the proposed SCP-based formulation computationally tractable while clarifying the intended modeling scope. The principal symbols used in the mathematical formulation are summarized in Appendix A.

2.1. UAV Point-Mass Dynamics and Linearization

In this study, a pseudo six-degree-of-freedom (6-DoF) point-mass model is adopted to describe the UAV motion for computational efficiency. The inertial reference frame is defined using the North–East–Down (NED) frame, and the state vector $X \in \mathbb{R}^7$ consists of the position $p = [p_N, p_E, p_D]^T$, the velocity $v = [v_N, v_E, v_D]^T$, and the yaw angle ψ . The control input $U \in \mathbb{R}^4$ comprises the thrust vector $T = [T_N, T_E, T_D]^T$ and the yaw rate ω . The state and control input are defined as follows:

$$X = [p, v, \psi]^T, \quad U = [T, \omega]^T. \quad (1)$$

In the dynamics model, the translational acceleration is defined as the resultant of the thrust vector T , gravity g , and aerodynamic drag D . The continuous-time dynamics are given as follows:

$$\dot{p} = v, \quad \dot{v} = \frac{1}{m} (T + D(v)) + g, \quad \dot{\psi} = \omega. \quad (2)$$

Here, m denotes the vehicle mass, and $g = [0, 0, g_0]^T$ is the gravitational acceleration vector expressed in the NED frame. In this study, the aerodynamic drag $D(v)$ is modeled as a quadratic force that acts opposite to the velocity direction with a magnitude proportional to the squared speed as follows:

$$D(v) = -\frac{1}{2} \rho C_D S_{\text{ref}} \|v\| v. \quad (3)$$

Here, ρ denotes the air density, C_D is the drag coefficient, and S_{ref} is the reference area. Since the drag model (3) is a nonlinear function of the velocity, it is linearized via a first-order Taylor expansion about the reference trajectory (\bar{X}_k, \bar{U}_k) at each SCP iteration. Specifically over the k -th

interval, we have that:

$$D(v_k) \approx \bar{D}_k + J_{v,k}^D (v_k - \bar{v}_k),$$

$$J_{v,k}^D = \left. \frac{\partial D}{\partial v} \right|_{v=\bar{v}_k} = -\frac{1}{2} \rho C_D S \left(\frac{\bar{v}_k \bar{v}_k^\top}{\|\bar{v}_k\|} + \|\bar{v}_k\| I_3 \right), \quad (4)$$

which yields a continuous-time linear affine model with respect to the state as follows:

$$\dot{X} \approx A_k X + B_k U + c_k, \quad (5)$$

with

$$A_k = \begin{bmatrix} 0_{3 \times 3} & I_3 & 0_{3 \times 1} \\ 0_{3 \times 3} & \frac{1}{m} J_{v,k}^D & 0_{3 \times 1} \\ 0_{1 \times 3} & 0_{1 \times 3} & 0 \end{bmatrix}, \quad B_k = \begin{bmatrix} 0_{3 \times 3} & 0_{3 \times 1} \\ \frac{1}{m} I_3 & 0_{3 \times 1} \\ 0_{1 \times 3} & 1 \end{bmatrix}, \quad (6)$$

$$c_k = \begin{bmatrix} 0_{3 \times 1} \\ \frac{1}{m} (\bar{D}_k - J_{v,k}^D \bar{v}_k) + g \\ 0 \end{bmatrix}.$$

Moreover, to address a free-final-time formulation in which the terminal time t_f is treated as a decision variable, we introduce a normalized time $\tau \in [0, 1]$ and define $t = t_f \tau$. Accordingly, since $dX/d\tau = t_f \dot{X}$, Equation (5) can be written in the normalized time domain as:

$$\frac{dX}{d\tau} \approx t_f (A_k X + B_k U + c_k). \quad (7)$$

In the SCP iterations, letting the reference terminal time be \bar{t}_f and decomposing $t_f = \bar{t}_f + \Delta t_f$, Equation (7) can be rearranged via a first-order approximation into the following form:

$$\frac{dX}{d\tau} \approx A_{\tau,k} X + B_{\tau,k} U + c_{\tau,k} + \Sigma_k \Delta t_f, \quad (8)$$

where

$$A_{\tau,k} = \bar{t}_f A_k, \quad B_{\tau,k} = \bar{t}_f B_k, \quad c_{\tau,k} = \bar{t}_f c_k, \quad (9)$$

$$\Sigma_k = A_k \bar{X}_k + B_k \bar{U}_k + c_k.$$

By explicitly including the first-order sensitivity term with respect to Δt_f , the free-final-time variable can be stably incorporated into the SCP framework in the form of linear constraints.

Finally, to discretize (8), this study adopts a matrix-exponential discretization based on the state transition matrix (STM). Assuming that $A_{\tau,k}$, $B_{\tau,k}$, $c_{\tau,k}$, and Σ_k are constant over the normalized time interval $[\tau_k, \tau_{k+1}]$ and letting $\Delta\tau = \tau_{k+1} - \tau_k = 1/N$, (8) can be interpreted as a linear time-invariant system on each interval and can be rearranged, including the input and constant terms, as

$$\frac{dX}{d\tau} = A_{\tau,k} X + B_{\tau,k} U + c_{\tau,k} + \Sigma_k \Delta t_f = A_{\tau,k} X + \hat{B}_k \hat{U}. \quad (10)$$

Here, we define the augmented input $\hat{U} \in \mathbb{R}^6$ and the corresponding input matrix \hat{B}_k as

$$\hat{U} = \begin{bmatrix} U \\ 1 \\ \Delta t_f \end{bmatrix}, \quad \hat{B}_k = [B_{\tau,k} \quad c_{\tau,k} \quad \Sigma_k]. \quad (11)$$




Figure 1: Definition of the UAV body-fixed coordinate frame and Euler angles.

To derive the matrix-exponential discretization, we introduce the augmented state $\hat{X} \in \mathbb{R}^{7+6}$ and obtain an equivalent linear dynamics model given by (12):

$$\hat{X} \triangleq \begin{bmatrix} X \\ \hat{U} \end{bmatrix}, \quad \frac{d\hat{X}}{d\tau} = \begin{bmatrix} A_{\tau,k} & \hat{B}_k \\ 0 & 0 \end{bmatrix} \hat{X} = \mathcal{M}_k \hat{X}. \quad (12)$$

Specifically, over $\Delta\tau$, the solution satisfies

$$\hat{X}(\tau_{k+1}) = \exp(\mathcal{M}_k \Delta\tau) \hat{X}(\tau_k). \quad (13)$$

Exploiting the block structure, the matrix exponential can be expressed as

$$\exp(\mathcal{M}_k \Delta\tau) = \begin{bmatrix} \Phi_k & \Gamma_k \\ 0 & I \end{bmatrix}, \quad (14)$$

where Φ_k is the state transition matrix and Γ_k is the discrete-time input matrix associated with the augmented input \hat{U} . Therefore, the discrete-time dynamics can be written as $X_{k+1} = \Phi_k X_k + \Gamma_k \hat{U}_k$. By decomposing \hat{U}_k and $\Gamma_k = [\Gamma_{u,k} \quad \Gamma_{c,k} \quad \Gamma_{s,k}]$, we finally obtain (15):

$$X_{k+1} = \Phi_k X_k + \Gamma_{u,k} U_k + \Gamma_{c,k} + \Gamma_{s,k} \Delta t_f. \quad (15)$$

2.2. Coordinate Frames and Attitude Kinematics

The UAV body-fixed coordinate frame (B) is defined as a 3–2–1 Euler rotation of the NED frame (I), with the body axes denoted by \hat{x} , \hat{y} , and \hat{z} . To emulate rotorcraft-like motion, the direction opposite to the thrust vector is defined as the \hat{z} -axis of the body frame. Once ψ is determined, the roll angle ϕ and pitch angle θ are geometrically obtained from the yaw-rotated axis $\hat{z}' = R(\psi)\hat{z}$. Specifically, let N' , E' , and D' denote the coordinate axes rotated by ψ from the inertial frame. Then, ϕ and θ can be derived such that the D' -axis aligns with the \hat{z}' -axis in this rotated frame as:

$$\theta = \text{atan2}(\hat{z}'_x, \hat{z}'_z),$$

$$\phi = -\text{atan2}(\hat{z}'_y, \sqrt{(\hat{z}'_x)^2 + (\hat{z}'_z)^2}). \quad (16)$$

Using the resulting attitude angles ϕ , θ , and ψ , the transformation matrix between the inertial frame and the body frame is constructed as $R(\psi, \theta, \phi) = R(\phi)R(\theta)R(\psi)$. The UAV body-fixed frame and the corresponding Euler-angle definition are summarized in Fig. 1.




Figure 2: Definition of the radar line-of-sight azimuth and elevation angles within the UAV body frame.




Figure 3: Inertial line-of-sight geometry and radar-relative angles for terrain exposure evaluation.

2.3. Attitude-Dependent Incidence Angles for RCS Evaluation

From the radar perspective, it is necessary to describe the relative attitude and detection geometry in order to quantify the detectability of the UAV. Accordingly, to define a radar-oriented attitude representation, the body-frame azimuth and elevation angles, Az and El , are introduced. Let the line-of-sight (LOS) vector between the UAV position p and the radar position p_{radar} be defined as $r = p - p_{\text{radar}}$. The LOS vector expressed in the body frame is computed as $r^B = R(\psi, \theta, \phi)r$, and Az and El are defined as the azimuth and elevation of r^B in the UAV body frame. The angles satisfy $Az \in [-\pi, \pi]$ and $El \in [0, \pi]$, and their geometric definition is illustrated in Fig. 2. These angles are used for radar SNR modeling and for training the neural network surrogate models.

2.4. Radar-Centric Line-of-Sight Geometry for Terrain Exposure

To avoid radar exposure, the UAV is required to perform NOE maneuvers behind terrain ridgelines. In this context, LOS exposure means that the straight radar-UAV viewing ray is not blocked by terrain, whereas terrain masking occurs

Table 1

Radar and environmental parameters for SNR calculation.

Symbol	Value	Unit	Description
λ	0.0375	m	Wavelength
P_t	6000	W	Peak power
G	39	dBi	Antenna gain
τ_p	0.0001	s	Pulse width
T_s	290	K	System temperature
L	0	dB	Loss
PRF	2200	Hz	Pulse repetition frequency

when this ray is intercepted by terrain before reaching the UAV. To determine whether the UAV is exposed to the radar, it is necessary to quantitatively describe the relative position geometry from the radar perspective. Accordingly, this study defines the radar-based LOS angles χ and γ as the azimuth and elevation of the radar LOS in the NED frame. The angles satisfy $\chi \in [-\pi, \pi]$ and $\gamma \in [-\pi/2, \pi/2]$, and the corresponding inertial LOS geometry is depicted in Fig. 3. These angles are used for training the neural network surrogate model for terrain masking.

2.5. Radar Signal-to-Noise Ratio (SNR) Modeling

In this study, the signal-to-noise ratio (SNR) is adopted as the detection metric to quantify the radar-UAV detectability. First, a precomputed UAV radar cross section (RCS) database, $\sigma_{\text{dB}}(Az, El)$, is utilized as a function of the radar incidence azimuth and elevation angles, Az and El . To use the RCS in a linear scale, it is converted as $\sigma_{\text{lin}}(Az, El) = 10^{\sigma_{\text{dB}}/10}$. In addition, the slant range d is defined to account for the radar-UAV separation, and the SNR is computed using a generalized single-pulse radar equation as a function of (Az, El, d) :

$$\text{SNR}(Az, El, d) = \frac{P_t G^2 \lambda^2 \sigma_{\text{lin}}(Az, El) \tau_p}{(4\pi)^3 d^4 k T_s L} F_{\text{eccl}}(d). \quad (17)$$

Here, $F_{\text{eccl}}(d)$ denotes the eclipsing factor arising from the transmit-receive switching process. It is incorporated to reflect the attenuation effect of the radar duty cycle determined by the pulse width τ and pulse repetition frequency (PRF), as well as the inter-pulse blind intervals, on the resulting SNR. The radar and environmental parameters used for the SNR calculation in (17) are summarized in Table 1. In (17), the gain and loss terms are used in linear scale. The tabulated values of G and L are specified in dBi and dB and are converted to linear quantities in the numerical implementation.

In this study, the SNR is computed using the `radareqsnr` function in the MATLAB Radar Toolbox, and is organized as a three-dimensional lookup table over $Az \in [-\pi, \pi]$, $El \in [0, \pi]$, and $d \in [100 \text{ m}, 60000 \text{ m}]$. An example slice of the computed SNR table at $d = 30,000 \text{ m}$ is shown in Fig. 4.




Figure 4: Radar signal-to-noise ratio (SNR) distribution across body-frame azimuth and elevation angles at a slant range of $d = 30,000$ m.

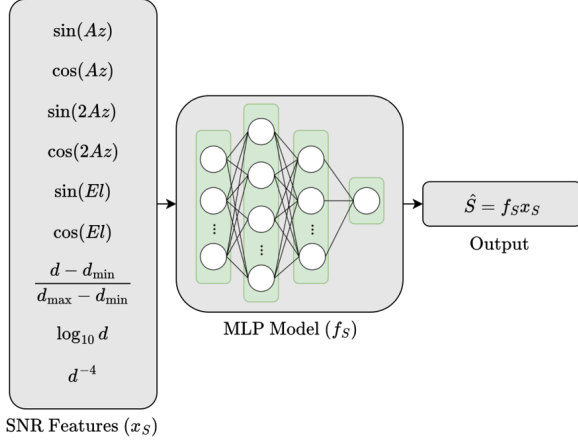


Figure 5: Architecture of the multilayer perceptron (MLP) neural surrogate model for SNR approximation.

3. Differentiable Surrogate Models for Survivability Metrics

3.1. Differentiable Neural Surrogate for SNR

The SNR computed in the previous step is a highly nonlinear quantity, resulting from the combined effects of the strong directional dependence of the RCS and range-induced attenuation. Therefore, directly using the precomputed table in trajectory optimization is not ideal, because interpolation over a highly nonlinear SNR map may yield non-smooth local behavior and makes it difficult to obtain reliable derivatives for SCP linearization. To address this, this study trains a multilayer perceptron (MLP)-based neural network and employs it as a continuous, smooth, and differentiable surrogate model to approximate the dominant SNR distribution patterns. The ground-truth SNR data are denoted by \mathcal{S} , and the input x_S , output $\hat{\mathcal{S}}$, and MLP model f_S are constructed as shown in Fig. 5, which summarizes the feature construction and the MLP mapping from x_S to $\hat{\mathcal{S}}$.

The inputs are preprocessed into learning-suitable features by using sin and cos terms to capture the periodic nature of Az and El . In addition, to reflect range attenuation and the characteristics of the radar equation, normalized and log-scaled range terms as well as a d^{-4} term are employed.

Table 2

Neural network architecture and training hyperparameters for the SNR surrogate model.

Item	Value
Architecture	MLP: 9–64–32–1
Activation	SiLU
Loss	Huber loss ($\delta = 0.5$)
Optimizer	AdamW
Learning rate	4×10^{-3}
Epochs	15000

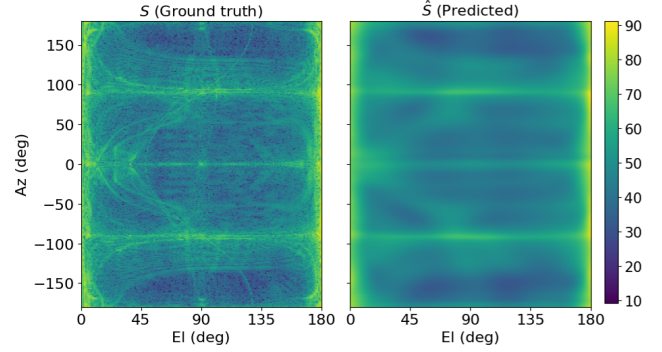


Figure 6: Comparison between the true and MLP-predicted SNR distributions at a slant range of $d = 30,000$ m.

The output is defined as the SNR data standardized via z-score normalization using the mean and standard deviation of the entire dataset, and it is converted back to the actual SNR in the dB scale after training for subsequent use. The SNR neural model f_S is designed with an MLP architecture; the SiLU activation is adopted to obtain smoother and differentiable nonlinearities compared to ReLU, and the loss function is chosen as the Huber loss $\ell_\delta(\epsilon_S)$ in (18). The Huber loss preserves a smooth and differentiable form similar to the mean squared error (MSE) when the residual ($\epsilon_S = \mathcal{S} - \hat{\mathcal{S}}$) is small, while it transitions to a mean absolute error (MAE)-like form for large residuals, thereby more strongly penalizing large errors. Here, δ denotes the threshold that separates the MSE and MAE regimes. Table 2 summarizes the architecture of the SNR neural surrogate model and the main training hyperparameters used in this study. Training is performed using the AdamW algorithm, and a representative comparison between the reconstructed \mathcal{S} and the predicted $\hat{\mathcal{S}}$ in the dB scale is shown in Fig. 6.

$$\ell_\delta(\epsilon_S) = \begin{cases} \frac{1}{2}\epsilon_S^2, & \text{if } |\epsilon_S| \leq \delta, \\ \delta(|\epsilon_S| - \frac{1}{2}\delta), & \text{otherwise.} \end{cases} \quad (18)$$

To characterize the surrogate fit on the training dataset, the MSE and R^2 metrics are computed, and the resulting values are reported in Table 3. For the trained model, the loss converges to approximately 0.1, as shown in Fig. 7. Because the ground-truth SNR surface contains sharp localized variations, the surrogate is not intended to reproduce every fine-scale peak exactly. Instead, it is used to provide a

Table 3
Approximation error metrics on the SNR training dataset.

Model	Performance	
	MSE	R^2
SNR neural surrogate	0.10966	0.89034




Figure 7: Training loss convergence curve for the SNR neural surrogate model.

differentiable approximation of the dominant SNR patterns required by the present optimization framework. This level of fit is therefore considered adequate for the present proof-of-concept study. In the present framework, the SNR surrogate is used as a smooth and differentiable proxy for the radar-equation-based SNR table within the SCP optimization process. Its primary role is to provide local gradients and capture the dominant attitude- and range-dependent SNR trends that guide the optimizer. The same surrogate output is also used as a consistent quantitative metric for comparing the simulated cases in Section 5. Therefore, the reported SNR values should be interpreted within the assumed radar equation, RCS database, and training domain used to construct the surrogate.

3.2. Differentiable Neural Surrogate for Terrain Masking

In mountainous terrain, the most critical factor for enhancing UAV survivability and minimizing detectability is to block visibility. To this end, it is essential for the UAV to be masked by terrain features with respect to the radar, and thus NOE flight at low altitude below terrain ridgelines is required.

In this study, we define an exposure index, denoted by E , to quantify the radar-UAV visibility exposure induced by terrain. To leverage E within the trajectory optimization process, a neural network surrogate model is similarly trained. For training, we use mountainous terrain data covering a 40 km \times 40 km area with a grid resolution of 25 m. The terrain is represented as a set of facets by subdividing each grid cell into two triangular facets. The inputs are the previously defined χ , γ , and d . Considering a typical engagement scenario in which the radar and the UAV are separated by tens of kilometers, the input domain is set to $\chi \in$

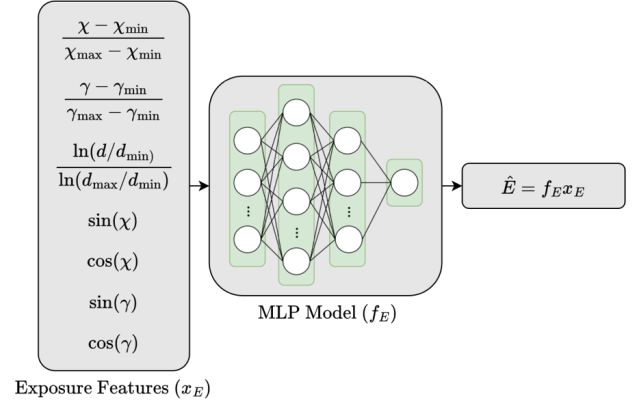


Figure 8: Architecture of the MLP neural surrogate model for the terrain exposure index (E).

Table 4
Neural network architecture and training hyperparameters for the terrain exposure surrogate model.

Item	Value
Architecture	MLP: 7–64–32–1
Activation	SiLU
Loss	Weighted Huber loss ($\delta = 0.1$)
Optimizer	AdamW
Learning rate	3×10^{-3}
Epochs	5000
Boundary weights	$W_{\max} = 1000$, $m_E = 0.2$

$[-\pi, \pi]$, $\gamma \in [-\pi/6, \pi/6]$, and $d \in [10000 \text{ m}, 40000 \text{ m}]$. These variables are converted into input features using log-scale and min–max normalization, together with sin and cos terms. The neural network ground-truth data are denoted by E , and the input features, output \hat{E} , and MLP-based neural surrogate model f_E are organized as shown in Fig. 8, which summarizes the feature construction and the mapping from the radar-relative geometry to the terrain exposure index.

The exposure index E is computed from the skyline elevation angle derived from the radar–terrain geometry. First, a unit LOS direction is constructed from (χ, γ) . Then, using the RayMeshIntersector in the Python trimesh library, we compute the first intersection distance between the terrain facets and the ray cast from the radar along the specified direction, denoted by $d_{\text{cr}}(\chi, \gamma)$. This corresponds to the nearest range at which the ray intersects the terrain. Based on this intersection, we define the skyline elevation angle γ_{sky} , i.e., the minimum elevation angle that guarantees an unobstructed line of sight beyond the slant range d , and the elevation-angle offset $\Delta\gamma$ as:

$$\begin{aligned} \gamma_{\text{sky}}(\chi, d) &= \inf \{ \gamma \mid d_{\text{cr}}(\chi, \gamma) > d \} \\ \Delta\gamma(\chi, \gamma, d) &= \gamma - \gamma_{\text{sky}}(\chi, d). \end{aligned} \quad (19)$$

Accordingly, the sign of $\Delta\gamma$ indicates whether the UAV is below or above the terrain skyline.

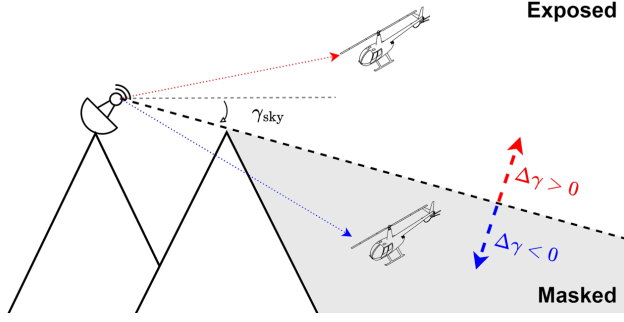


Figure 9: Geometric interpretation of the elevation-angle offset $\Delta\gamma$ for exposed and terrain-masked conditions.

Fig. 9 illustrates the geometric meaning of $\Delta\gamma$. The quantity $\Delta\gamma$ is an elevation-angle offset measured relative to the terrain skyline. Therefore, its sign directly determines the visibility regime: $\Delta\gamma > 0$ corresponds to a radar-exposed condition above the skyline, $\Delta\gamma = 0$ corresponds to the visibility transition boundary, and $\Delta\gamma < 0$ corresponds to a terrain-masked condition below the skyline. Its magnitude represents the angular separation from the skyline. A larger positive value indicates stronger geometric exposure, whereas a more negative value indicates greater separation below the skyline.

To reflect the characteristics that the geometric exposure level increases with larger $\Delta\gamma$, and that the masking separation increases as the UAV lies farther below the skyline for $\Delta\gamma < 0$, we construct the exposure index E using a Softplus-based hinge function ζ that is monotonically increasing in $\Delta\gamma$ and saturates to a lower bound of -1 . With this design, E robustly maintains a negative saturated value in the occluded region, while it increases with $\Delta\gamma$ in the exposed region. The resulting exposure index E is defined as:

$$\zeta(x) = \eta \ln(1 + e^{x/\eta}),$$

$$E(\chi, \gamma, d) = -1 + \frac{\zeta(\kappa \Delta\gamma(\chi, \gamma, d) + \nu)}{\zeta(\nu)}. \quad (20)$$

Here, η , κ , and ν are constants that determine the smoothing level, slope scaling, and saturation level, respectively, and are numerically tuned in this study. Accordingly, the sign of E preserves the same visibility interpretation in the optimization: negative values indicate terrain-masked flight, values near zero indicate the visibility transition boundary, and positive values indicate radar-exposed flight.

The model f_E adopts an MLP architecture, consistent with the SNR neural surrogate, and employs the SiLU activation and the AdamW optimizer. As the loss function, we use a Huber-loss-based weighted objective. In particular, accurate learning near the visibility transition boundary ($E \approx 0$) is crucial. Therefore, a Gaussian-kernel weighting function $w(E)$ is applied so that samples with E closer to zero contribute more strongly to the loss, and the corresponding




Figure 10: Comparison between the true and MLP-predicted terrain exposure index (E) at a slant range of $d = 30,000$ m.




Figure 11: Visual evaluation of the visibility transition boundary ($E \approx 0$) approximation at a slant range of $d = 30,000$ m.

weighted Huber loss ℓ_E is defined as:

$$w(E) = 1 + (W_{\max} - 1) \exp(-E/m_E)^2),$$

$$\ell_E = \frac{1}{M} \sum_{i=1}^M w(E_i) \ell_\delta(\epsilon_{E,i}). \quad (21)$$

Here, M denotes the number of samples in the training dataset. In the weighting function $w(E)$, W_{\max} denotes the maximum weight, and m_E is a parameter that determines the decay width of the weight away from the boundary. The residual is defined as $\epsilon_E = E - \hat{E}$. Table 4 summarizes the architecture of the terrain exposure neural surrogate model and the main training hyperparameters used in this study.

The learned terrain exposure surrogate is evaluated by comparing the ground-truth E and the prediction \hat{E} , as shown in Fig. 10. Moreover, to specifically assess the approximation quality around the visibility transition boundary, Fig. 11 provides a focused visualization of the region near $E \approx 0$ at a representative slant range of $d = 30,000$ m.

To characterize the surrogate fit on the training dataset, we compute the MSE and R^2 metrics. The results are reported in Table 5, and the same metrics are additionally computed for the boundary subset satisfying $|E| < 0.2$ to assess approximation quality near the visibility transition boundary.

The training process exhibits stable convergence, as illustrated by the loss curve in Fig. 12. In the present framework, the terrain exposure surrogate is used as a differentiable proxy for the LOS-based exposure calculation over the specified terrain map and input domain. Its output provides the exposure-dependent switching information used in the

Table 5

Approximation error metrics on the terrain exposure training dataset, including the boundary subset $|E| < 0.2$.

Data set	Performance	
	MSE	R^2
All data	0.00147	0.99983
Boundary region ($ E < 0.2$)	0.00082	0.94199




Figure 12: Training loss convergence curve for the terrain exposure neural surrogate model.

SCP formulation and is also used to compute the exposure ratio in the simulation analysis. Accordingly, the exposure predictions and ER values are interpreted with respect to the terrain geometry, radar location, and sampling ranges considered in this study.

3.3. CVaR-Based Conservative Margin for Visibility Uncertainty

The exposure index E designed above exhibits false-negative (FN) cases near the visibility transition boundary ($|E| < 0.2$), where approximation errors cause truly exposed states to be misclassified as occluded. To quantify this effect, we extract boundary-adjacent samples and analyze the sign-mismatch rate in this region. In particular, the most consequential misclassification occurs when $E \geq 0$ but the model predicts $\hat{E} < 0$, which leads to an FN event. Over the boundary-adjacent samples, the measured FN rate is 6.93351%. Because the optimizer may operate near skyline boundaries, a conservative mitigation of such FN errors is introduced.

To this end, we define an FN violation measure on the boundary-adjacent region and introduce a conditional value-at-risk (CVaR)-based conservative safety margin m_{FN} that accounts for the tail risk of FN violations [15]. First, we define the candidate sample set near the exposure boundary as

$$\mathcal{D} = \{i \mid 0 \leq E \leq \Delta\}. \quad (22)$$

Here, i denotes the sample index at (χ_i, γ_i, d_i) . To measure the severity of FN misclassification over \mathcal{D} , we define the FN violation magnitude v_i as

$$v_i = [-\hat{E}_i]_+ \quad \text{for } i \in \mathcal{D}. \quad (23)$$

Table 6

False-negative (FN) rates and the CVaR-based conservative margin near the visibility transition boundary.

Metric (boundary region)	Value
Boundary definition Δ	0.2
FN rate (before margin)	6.93351%
FN rate (after margin)	0.65028%
Confidence level α	0.95
$\text{CVaR}_\alpha(v)$	0.04031

That is, $v_i > 0$ only when a truly boundary-exposed sample ($E \geq 0$) is predicted as occluded ($\hat{E}_i < 0$), thereby quantifying the magnitude of FN-induced occlusion. Conversely, if $\hat{E}_i \geq 0$, then $v_i = 0$.

To obtain a conservative margin against worst-case FN violations, we employ the CVaR at confidence level $\alpha \in (0, 1)$, defined as

$$\text{CVaR}_\alpha(v) = \min_{z \in \mathbb{R}} \left[z + \frac{1}{1 - \alpha} \mathbb{E} \left[(v - z)_+ \right] \right], \quad (24)$$

where z is an auxiliary variable and $(v - z)_+$ accumulates only the portion of v exceeding z . In this study, we set $\alpha = 0.95$ to reflect the worst 5% tail of FN violations and define the safety margin as

$$m_{\text{FN}} \triangleq \text{CVaR}_\alpha(\{v_i\}_{i \in \mathcal{D}}). \quad (25)$$

In this study, the boundary region is defined with $\Delta = 0.2$, and the CVaR at confidence level $\alpha = 0.95$ is computed to determine a conservative margin of $m_{\text{FN}} = 0.04031$. The resulting CVaR-based safety margin is incorporated into the exposure prediction \hat{E}_k by defining

$$\hat{E}^{\text{rob}} \triangleq \hat{E} + m_{\text{FN}}. \quad (26)$$

By adding m_{FN} , FN misclassifications are conservatively suppressed, thereby reducing the risk of incorrectly classifying truly exposed intervals as terrain-masked and deactivating the exposed-regime SNR reduction term. To assess the effect of the margin on the boundary-adjacent samples used in the margin design, we re-evaluate the FN rate and observe a reduction from 6.93351% (without the margin) to 0.65028% (with the margin), as summarized in Table 6. This result indicates that the proposed margin provides a conservative treatment of boundary-adjacent FN errors and stabilizes the subsequent exposure/occlusion switching logic near ridgeline boundaries, thereby affecting which objective terms are activated near the visibility transition boundary.

3.4. Dynamic Weapon Engagement Zone (WEZ) Modeling

As introduced in Section 1, a weapon engagement zone (WEZ) denotes the region in which a SAM system can effectively engage the UAV. In the trajectory optimization problem, this region is treated as a time-varying threat constraint rather than as a post-processing safety check. This treatment is particularly important for rotorcraft-like NOE




Figure 13: Geometry of the cardioid-shaped horizontal weapon engagement zone.

flight, because low-altitude operation near terrain reduces the available clearance and maneuvering margin after a missile launch. Therefore, the proposed formulation incorporates WEZ avoidance directly as a state constraint, allowing the UAV to proactively maintain separation from the SAM threat region.

In this study, we assume a cardioid-shaped horizontal engagement radius whose extent depends on the pointing direction of the SAM launcher [8]. Accordingly, the engagement range is shaped by the launcher pointing direction, so that the effective WEZ is reduced in directions away from the SAM launcher orientation. This cardioid-shaped model is adopted instead of a circular WEZ because the effective engagement capability of a SAM is not isotropic with respect to the UAV motion. A circular WEZ assumes the same engagement range in all directions and therefore cannot capture the dependence of missile reachability on the relative LOS geometry and the UAV direction of motion. In contrast, the cardioid-shaped WEZ varies its size according to the relative angle between the UAV horizontal velocity and the SAM–UAV LOS direction. As a result, the threat boundary expands when the UAV approaches under a favorable engagement geometry for the SAM, and contracts when the UAV motion induces a less favorable engagement geometry. This provides a more geometry-aware and less conservative threat representation than a fixed circular boundary, while retaining a simple analytical form suitable for trajectory optimization. Let θ_h denote the launcher pointing direction, χ_h the horizontal LOS bearing from the SAM to the UAV, and ξ the relative angle between the UAV horizontal velocity direction and the horizontal LOS direction. Then, the WEZ radius is modeled as:

$$\rho(\theta_h, \chi_h, \xi) = \frac{R_{\max}}{2} \frac{(1 + \cos \xi)}{2} \left(1 + \sin\left(\frac{\pi}{2} - \chi_h + \theta_h\right) \right),$$

where R_{\max} denotes the maximum horizontal engagement range of the SAM. Fig. 13 illustrates the cardioid-shaped horizontal WEZ geometry.

The WEZ radius becomes largest when the launcher pointing direction aligns with the LOS direction, i.e., $\theta_h =$

χ_h , which corresponds to the worst-case condition in which the engagement capability is maximized toward the UAV. Therefore, in the optimization problem, the launcher pointing direction is treated through this worst-case alignment rather than by prescribing a specific value of θ_h . In addition, the term $1 + \cos \xi$ is maximized as the UAV approaches along the LOS direction. The relative angle $\xi_{j,k}$ is defined via the inner product between the UAV horizontal unit velocity vector $\hat{v}_{h,k} \in \mathbb{R}^2$ at time step k and the horizontal LOS unit vector $\hat{r}_{h,j,k} \in \mathbb{R}^2$ pointing from the UAV toward SAM j :

$$\begin{aligned} \cos \xi_{j,k} &= \hat{v}_{h,k}^\top \hat{r}_{h,j,k}, & \hat{v}_{h,k} &= \frac{v_{h,k}}{\|v_{h,k}\|}, \\ \hat{r}_{h,j,k} &= \frac{p_{\text{SAM},j} - p_{h,k}}{\|p_{\text{SAM},j} - p_{h,k}\|}, \end{aligned} \quad (27)$$

where $v_{h,k} \in \mathbb{R}^2$ is the UAV horizontal velocity at time step k , and $p_{\text{SAM},j} \in \mathbb{R}^2$ and $p_{h,k} \in \mathbb{R}^2$ denote the horizontal positions of SAM j and the UAV at time step k , respectively. Under this worst-case launcher alignment, the direction-dependent WEZ model reduces to the following maximum-radius form:

$$\rho_{\max}(\xi) = \frac{R_{\max}}{2} (1 + \cos \xi). \quad (28)$$

Defining the horizontal separation as $d_{h,j,k} = \|p_{h,k} - p_{\text{SAM},j}\|_2$, WEZ avoidance can be enforced by requiring the UAV to remain outside the maximum-radius boundary at every time step, which corresponds to avoiding the worst-case envelope of the direction-dependent WEZ model. This yields the following constraint used in the subsequent optimization problem. Here, \mathcal{A} denotes the set of SAM systems:

$$d_{h,j,k} \geq \rho_{\max}(\xi_{j,k}), \quad \text{for } j \in \mathcal{A}. \quad (29)$$

4. SCP-Based Trajectory Optimization Formulation

4.1. SDF-Based Terrain Collision Avoidance

In this study, a three-dimensional terrain model is constructed from terrain data using triangular facets. First, each cell of the terrain grid is subdivided into two triangular facets, and the index set of all facets comprising the terrain is denoted by \mathcal{O} . For each facet $i \in \mathcal{O}$, the facet centroid $p_{\text{facet},i} \in \mathbb{R}^3$ is computed as the average of its three vertices, and the corresponding unit normal vector $n_i \in \mathbb{R}^3$ is extracted. The normal direction n_i is chosen to point from the terrain surface toward the free space, so that the n_i direction represents the outward direction away from the terrain. The quantities $p_{\text{facet},i}$ and n_i are computed using the Python trimesh library. Fig. 14 illustrates the triangular-facet decomposition of the 40 km \times 40 km mountainous terrain model used in this study.

For terrain collision checking, we employ a signed distance function (SDF) [16]-based surface representation. The SDF is a scalar function that returns the minimum signed




Figure 14: Triangular-facet representation of the 40 km×40 km mountainous terrain.

distance from an arbitrary point $q \in \mathbb{R}^3$ to the surface $\partial\Omega$. Its mathematical definition is given using

$$\text{dist}(q, \partial\Omega) = \min_{b \in \partial\Omega} \|q - b\|, \quad (30)$$

by

$$\text{SDF}(q) = \begin{cases} \text{dist}(q, \partial\Omega), & \text{if } q \in \Omega^C \\ 0, & \text{if } q \in \partial\Omega \\ -\text{dist}(q, \partial\Omega), & \text{if } q \in \Omega \end{cases} \quad (31)$$

where Ω denotes the set occupied by the object, Ω^C denotes the exterior region, and $\text{dist}(q, \partial\Omega)$ is the minimum Euclidean distance to the surface. Accordingly, the SDF returns a positive value when q lies outside the object (no collision), zero when q lies on the surface, and a negative value when q lies inside the object (collision).

To implement terrain collision avoidance constraints in a linear form, we approximate the SDF locally using a first-order planar model and formulate a half-space separation condition. Given a facet centroid $p_{\text{facet},i}$ and its outward unit normal n_i , the signed distance at the UAV position p_k can be approximated by the corresponding supporting half-space as:

$$\text{SDF}(p_k) = n_i^\top (p_k - p_{\text{facet},i}), \quad \text{for } i \in \mathcal{O}_k, \quad (32)$$

which is consistent with the outward half-space expression $n_i^\top (p - p_{\text{facet},i}) \geq 0$. Using this approximation, terrain collision avoidance is enforced by requiring a minimum clearance distance z_{safe} from the terrain surface:

$$n_i^\top (p_k - p_{\text{facet},i}) \geq z_{\text{safe}}, \quad \text{for } i \in \mathcal{O}_k. \quad (33)$$

Here, \mathcal{O}_k denotes the index set of facets adjacent to the UAV position p_k at time step k , and the constraint is imposed only over the selected facets in \mathcal{O}_k . Fig. 15 illustrates the local half-space interpretation of the terrain clearance constraint and the corresponding feasible region defined by z_{safe} .

4.2. Trajectory Optimization via Sequential Convex Programming (SCP)

Directly embedding the neural surrogate outputs \hat{S}_k and \hat{E}_k into the trajectory optimization problem yields a highly nonconvex formulation due to the coupled nonlinear




Figure 15: Illustration of the local half-space terrain clearance constraint based on the facet centroid and outward normal.

flight dynamics, attitude-dependent radar geometry, terrain-related visibility changes, and dynamic WEZ constraints. To obtain a computationally tractable solution, this study adopts sequential convex programming (SCP), in which a sequence of locally valid convex subproblems is iteratively constructed around a reference trajectory.

The tractability of the proposed framework comes from the distinct computational roles of the surrogate models, automatic differentiation, and SCP. The surrogate models provide smooth and differentiable proxy functions for complex survivability metrics that are difficult to embed directly in the optimization problem. The SNR surrogate avoids relying on direct interpolation of the highly nonlinear radar-equation-based SNR table and provides differentiable outputs suitable for SCP linearization. The terrain exposure surrogate similarly replaces repeated LOS-based terrain-visibility calculations over the terrain geometry with a differentiable approximation over the specified input domain. Automatic differentiation is used to compute the Jacobians of the composed mappings involving UAV position, thrust, yaw, radar-angle geometry, radar range, and neural surrogate outputs, avoiding cumbersome manual derivation and repeated finite-difference perturbations. SCP then avoids solving the full nonlinear and nonconvex trajectory problem directly by replacing the dynamics and other nonlinear quantities, including the surrogate outputs and WEZ radius, with first-order affine approximations around the current reference trajectory. As a result, each iteration is formulated as a convex subproblem, followed by reference update and trust-region adjustment.

At each SCP iteration, the nonlinear quantities required by the detectability and terrain masking models are replaced by first-order affine approximations about the current reference trajectory (\bar{X}_k, \bar{U}_k) . First, the body-frame radar azimuth and elevation angles are linearized as

$$\begin{bmatrix} Az_k \\ El_k \end{bmatrix} \approx \begin{bmatrix} \bar{A}z_k \\ \bar{E}l_k \end{bmatrix} + J_{p,k}^{AE} (p_k - \bar{p}_k) + J_{T,k}^{AE} (T_k - \bar{T}_k) + J_{\psi,k}^{AE} (\psi_k - \bar{\psi}_k), \quad (34)$$

where $J_{p,k}^{AE}$, $J_{T,k}^{AE}$, and $J_{\psi,k}^{AE}$ denote the Jacobians with respect to position, thrust, and yaw, respectively. Similarly,

the radar-centric LOS angles used for terrain masking are linearized as

$$\begin{bmatrix} \chi_k \\ \gamma_k \end{bmatrix} \approx \begin{bmatrix} \bar{\chi}_k \\ \bar{\gamma}_k \end{bmatrix} + J_{p,k}^{\chi\gamma} (p_k - \bar{p}_k), \quad (35)$$

where $J_{p,k}^{\chi\gamma}$ is the Jacobian of the mapping from p_k to (χ_k, γ_k) .

Using these affine geometric quantities, the surrogate outputs are linearized with respect to their inputs. The SNR surrogate output is approximated by

$$\begin{aligned} \hat{S}_k \approx \bar{S}_k + \frac{\partial \hat{S}_k}{\partial Az_k} (Az_k - \bar{A}z_k) + \frac{\partial \hat{S}_k}{\partial El_k} (El_k - \bar{E}l_k) \\ + \frac{\partial \hat{S}_k}{\partial d_k} (d_k - \bar{d}_k), \end{aligned} \quad (36)$$

and the terrain masking surrogate output is approximated as

$$\hat{E}_k \approx \bar{E}_k + \frac{\partial \hat{E}_k}{\partial \chi_k} (\chi_k - \bar{\chi}_k) + \frac{\partial \hat{E}_k}{\partial \gamma_k} (\gamma_k - \bar{\gamma}_k) + \frac{\partial \hat{E}_k}{\partial d_k} (d_k - \bar{d}_k). \quad (37)$$

The dynamic WEZ model is treated in the same manner. Since the maximum engagement radius in (29) depends nonlinearly on the horizontal relative geometry through $\xi_{j,k}$, it is written as $\rho_{\max,j}(p_{h,k}, v_{h,k})$ and linearized around the reference trajectory as

$$\begin{aligned} \rho_{\max,j}(p_{h,k}, v_{h,k}) \approx \bar{\rho}_{\max,j,k} + J_{p,j,k}^\rho (p_{h,k} - \bar{p}_{h,k}) \\ + J_{v,j,k}^\rho (v_{h,k} - \bar{v}_{h,k}), \end{aligned} \quad (38)$$

where $J_{p,j,k}^\rho$ and $J_{v,j,k}^\rho$ denote the Jacobians with respect to the horizontal position and velocity, respectively.

In the implementation, the Jacobians associated with the radar-angle mappings and the neural surrogate models in (34)–(37) are evaluated using PyTorch automatic differentiation, whereas nonlinear quantities with explicit closed-form expressions, such as the drag Jacobian and the WEZ radius Jacobian in (38), are computed analytically. Other nonlinear quantities appearing in the SCP subproblem, including the roll and pitch angles reconstructed from the thrust vector, the radar slant range d_k , the SAM-relative horizontal distance $d_{h,j,k}$, the desired heading angle $\psi_{\text{des},k} = \text{atan2}(v_{E,k}, v_{N,k})$, and the body-axis directional speed components, are linearized in the same first-order manner. Here, $\psi_{\text{des},k}$ is introduced to align the UAV heading with the ground-projected velocity direction during terrain-masked flight.

To account for regime-dependent maneuvering according to the radar visibility condition, the robust terrain masking metric \hat{E}_k^{rob} defined in (26) is used to construct the switching functions

$$\begin{aligned} \sigma_k^E &= [-\hat{E}_k^{\text{rob}}]_+, \\ \sigma_k^S &= [\hat{E}_k^{\text{rob}}]_+. \end{aligned} \quad (39)$$

Here, negative values are clipped to zero so that each maneuvering component is activated only in the appropriate regime. Specifically, in terrain-masked segments, σ_k^E

activates the heading-alignment term, whereas in exposed segments, σ_k^S activates the SNR reduction term. Because these coefficients are computed from \hat{E}_k^{rob} , which includes the CVaR margin, the switching functions influence the optimization by shifting the exposed/masked switching boundary and, consequently, the time intervals over which the heading-alignment and SNR reduction terms are active. In each SCP subproblem, these switching coefficients are evaluated from the previous reference trajectory and then held fixed during the convex optimization step in order to preserve convexity.

In addition, trust-region constraints are imposed at every SCP iteration so that the update remains within a local neighborhood where the first-order affine approximations are reliable. To this end, separate trust-region radii are assigned to the terminal time, position, velocity, and yaw angle, denoted by Δ_{t_f} , Δ_p , Δ_v , and Δ_ψ , respectively. The corresponding bounds are given by

$$\begin{aligned} |t_f - \bar{t}_f| \leq \Delta_{t_f}, \quad \|p_k - \bar{p}_k\|_2 \leq \Delta_p, \\ \|v_k - \bar{v}_k\|_2 \leq \Delta_v, \quad |\psi_k - \bar{\psi}_k| \leq \Delta_\psi. \end{aligned} \quad (40)$$

Here, Δ_{t_f} limits the update of the terminal time, while Δ_p , Δ_v , and Δ_ψ bound the allowable deviations of the position, velocity, and yaw angle from the reference trajectory, respectively. These bounds prevent excessively large steps that may invalidate the local convex approximation and thereby improve numerical stability. Furthermore, the trust-region radii are adjusted at each iteration using both the ratio between the actual nonlinear cost reduction and the predicted reduction obtained from the convex subproblem, together with the measured linearization error.

Using the predicted SNR \hat{S}_k and the terrain exposure index \hat{E}_k , the trajectory optimization problem is formulated to reduce the surrogate-based detection metric while maintaining feasible flight and satisfying terrain and threat constraints. The objective function consists of a terminal time minimization term, a maneuvering-effort term minimizing \dot{v}_k , a yaw rate regularization term minimizing ω_k , a terrain masking term promoting low-exposure flight, and an SNR reduction term that becomes active in exposed regions. In this way, the optimized trajectory is encouraged to exploit terrain masking whenever possible, while performing attitude-coupled detectability reduction maneuvers when exposure is unavoidable.

The constraints are imposed as follows: the discretized dynamics constraint, initial and terminal boundary conditions, maximum thrust magnitude and thrust-direction constraints, terrain collision avoidance constraints, longitudinal/lateral/climb/descent speed limits, dynamic WEZ avoidance constraints, the trust-region bounds, and a heading-alignment term during terrain-masked flight. The directional speed limits are introduced to reflect rotorcraft maneuvering limitations along each axis [17]. Accordingly, the simulation study should be interpreted as representing a rotorcraft-type UAV or VTOL platform operating in a low-altitude NOE mission profile, rather than a generic fixed-wing vehicle.

Table 7

Baseline simulation parameters, mission geometry, and UAV dynamic constraints.

Symbol	Value	Unit	Description
N	500	-	Number of discretization intervals
m	10^4	kg	UAV mass
g_0	9.81	m/s^2	Gravitational acceleration
T_{\max}	$2mg_0$	N	Maximum thrust
$ \mathcal{O}_k $	4	-	Number of facets adjacent to the UAV at time step k
z_{safe}	5	m	Minimum separation distance
p_{init}	$[600, 44100, -500]^\top$	m	Initial position of UAV in the NED frame
p_{fin}	$[36000, 9800, -700]^\top$	m	Terminal position of UAV in the NED frame
p_{radar}	$[3700, 15300, -1200]^\top$	m	Radar position in the NED frame
$p_{\text{SAM},1}$	$[16500, 31700]^\top$	m	Horizontal position of SAM site 1
$p_{\text{SAM},2}$	$[34800, 18400]^\top$	m	Horizontal position of SAM site 2
$ \mathcal{A} $	2	-	Number of SAM sites
R_{\max}	13000 / 12000	m	Maximum range of WEZ
ψ_{tol}	0.1	deg	Maximum allowable yaw error
\bar{v}_{fwd}	77.17	m/s	Maximum forward speed
\bar{v}_{rear}	23.15	m/s	Maximum rearward speed
\bar{v}_{lat}	23.15	m/s	Maximum lateral speed
\bar{v}_{climb}	12.7	m/s	Maximum climb rate
\bar{v}_{descent}	15.24	m/s	Maximum descent rate

We finally present the convex SCP subproblem that we solve in the following:

$$\begin{aligned}
 \text{minimize} \quad & \lambda_{t_f} t_f + \lambda_U \sum_{k=0}^{N-1} \|\dot{v}_k\|^2 + \lambda_\omega \sum_{k=0}^{N-1} \omega_k^2 \\
 & + \lambda_E \sum_{k=0}^{N-1} \hat{E}_k + \lambda_S \sum_{k=0}^{N-1} \sigma_k^S \hat{S}_k^2 \\
 \text{subject to} \quad & (34), (35), (36), (37), (38), \\
 & (15), (40), \\
 & p_0 = p_{\text{init}}, \quad p_N = p_{\text{fin}}, \\
 & \|T_k\| \leq T_{\max}, \quad T_{D,k} \leq 0, \\
 & -v_{\text{lat}} \leq v_{\text{lat},k} \leq v_{\text{lat}}, \\
 & -v_{\text{rear}} \leq v_{\text{long},k} \leq v_{\text{fwd}}, \\
 & -v_{\text{climb}} \leq v_{\text{vert},k} \leq v_{\text{descent}}, \\
 & n_i^\top (p_k - p_{\text{facet},i}) \geq z_{\text{safe}}, \quad \forall i \in \mathcal{O}_k, \\
 & d_{h,j,k} \geq \rho_{\max,j}(p_{h,k}, v_{h,k}), \quad \forall j \in \mathcal{A}, \\
 & \sigma_k^E (|\psi_k - \psi_{\text{des},k}| - \psi_{\text{tol}}) \leq 0.
 \end{aligned} \tag{41}$$

Here, the state variables are defined at all nodes $k = 0, \dots, N$, while the control inputs and other stage-wise quantities are defined over the intervals $k = 0, \dots, N-1$. Each convex SCP subproblem is solved using the GUROBI optimizer. To clarify the implementation sequence, a simplified pseudocode of the overall computational procedure is provided in Appendix B.

5. Simulation Results and Discussion


In this study, we use the same mountainous terrain dataset employed for training the terrain masking neural surrogate model. The terrain is converted into a triangular-facet representation and is used for both the terrain collision avoidance constraints and LOS-based occlusion evaluation. The radar location p_{radar} is placed at a mountain peak within the terrain, and the UAV starts from the initial position p_{init} , avoids two WEZ regions centered at $p_{\text{SAM},1}$ and $p_{\text{SAM},2}$, and reaches the terminal position p_{fin} . The mission geometry, common simulation parameters, and UAV dynamic constraints used throughout the simulations are summarized in Table 7.

Prior to the main simulations, we analyze the influence of the terrain masking term and the SNR reduction term on the resulting solutions. Fig. 16 shows how the exposure ratio ER and the mean surrogate-predicted SNR during exposed segments \bar{S}_E vary as the terrain masking weight λ_E and the SNR weight λ_S are adjusted. Here, ER is defined as the fraction of flight intervals satisfying the robust exposure criterion $\hat{E}^{\text{rob}} \geq 0$, and \bar{S}_E is computed by averaging the surrogate-predicted SNR over those exposed intervals. In each experiment, only one weight is varied while the remaining weights are kept fixed, and a total of 20 simulations are conducted. The horizontal axis in Fig. 16 reports the corresponding weight values normalized to $[0, 1]$ via linear scaling, denoted by $\tilde{\lambda}_E$ and $\tilde{\lambda}_S$. As $\tilde{\lambda}_E$ increases, the overall exposure ratio ER generally decreases, indicating that a larger terrain masking weight more strongly promotes occluded NOE-like flight. Likewise, increasing $\tilde{\lambda}_S$ reduces the mean exposed-segment surrogate-predicted SNR \bar{S}_E , suggesting that the SNR reduction term effectively lowers the predicted detection metric during unavoidable exposure.

Table 8

Summary of the simulation cases.

Case	Terrain masking	SNR reduction	Heading alignment	Role
Case 1	Off ($\tilde{\lambda}_E = 0$)	Off ($\tilde{\lambda}_S = 0$)	Always active	Baseline without survivability terms
Case 2	On ($\tilde{\lambda}_E = 0.33$)	Off ($\tilde{\lambda}_S = 0$)	Always active	Path-only terrain masking baseline
Case 3	On ($\tilde{\lambda}_E = 0.33$)	On ($\tilde{\lambda}_S = 1$)	Active in masked segments	Joint path-attitude comparison case
Case 4	On ($\tilde{\lambda}_E = 1$)	On ($\tilde{\lambda}_S = 1$)	Active in masked segments	Joint path-attitude case with stronger terrain masking

**Figure 16:** Weight-sweep results for terrain masking and SNR reduction: exposure ratio and mean exposed-segment SNR versus normalized weights.

The simulation cases and their purposes are summarized in Table 8. Cases 1 and 2 are separately generated baseline cases, whereas Cases 3 and 4 are selected from the weight-sweep study. Case 1 is a basic baseline in which neither terrain masking nor SNR reduction is included, and therefore represents a trajectory without survivability-oriented weighting. Case 2 is introduced as a path-only terrain masking baseline, where the terrain masking term is activated, the SNR reduction term is deactivated, and the heading-alignment constraint is imposed over the entire trajectory. Thus, Case 2 can optimize the geometric flight path to exploit terrain masking, but it cannot intentionally select lower-SNR attitudes during exposed segments. Case 3 uses the same terrain masking weight as Case 2, but activates the SNR reduction term and applies heading alignment in terrain-masked segments. Therefore, the comparison between Cases 2 and 3 evaluates the benefit of adding attitude optimization to path optimization under comparable terrain masking conditions. Case 4 represents a joint path-attitude case with a stronger terrain masking preference.

The quantitative metrics for the four cases are summarized in Table 9. The table reports ER, \tilde{S}_E , trajectory length, final flight time t_f , and the minimum WEZ margin using the definitions introduced above. The trajectory length is evaluated by summing the three-dimensional distances between consecutive trajectory nodes. The minimum WEZ margin is defined as the minimum value of $d_{h,j,k} - \rho_{\max,j,k}$

over all SAM sites and time steps; therefore, a positive value indicates that the WEZ avoidance constraint is satisfied throughout the trajectory.

The results for Case 1 are shown in Fig. 17 and Table 9. In this baseline case, the UAV heading follows ψ_{des} throughout the entire trajectory, and neither terrain masking nor SNR reduction is explicitly encouraged. Accordingly, the state and attitude histories do not show additional maneuvers for survivability enhancement. Quantitatively, Case 1 gives ER = 90.4% and $\tilde{S}_E = 55.70$ dB, which serve as reference values for evaluating the effect of the survivability-oriented terms.

Case 2, shown in Fig. 18, represents the path-only terrain masking baseline. Because the terrain masking term is activated, the geometric path is modified to reduce radar exposure. However, since the SNR reduction term is not included and heading alignment is imposed over the entire trajectory, ψ remains aligned with ψ_{des} even during exposed segments. As a result, the Az and El histories do not show intentional attitude steering toward low-SNR regions. Table 9 reports ER = 36.67% and $\tilde{S}_E = 52.96$ dB for this case. These results indicate that path optimization alone can affect the exposure characteristics through terrain masking, but it does not directly exploit attitude-dependent radar observability.

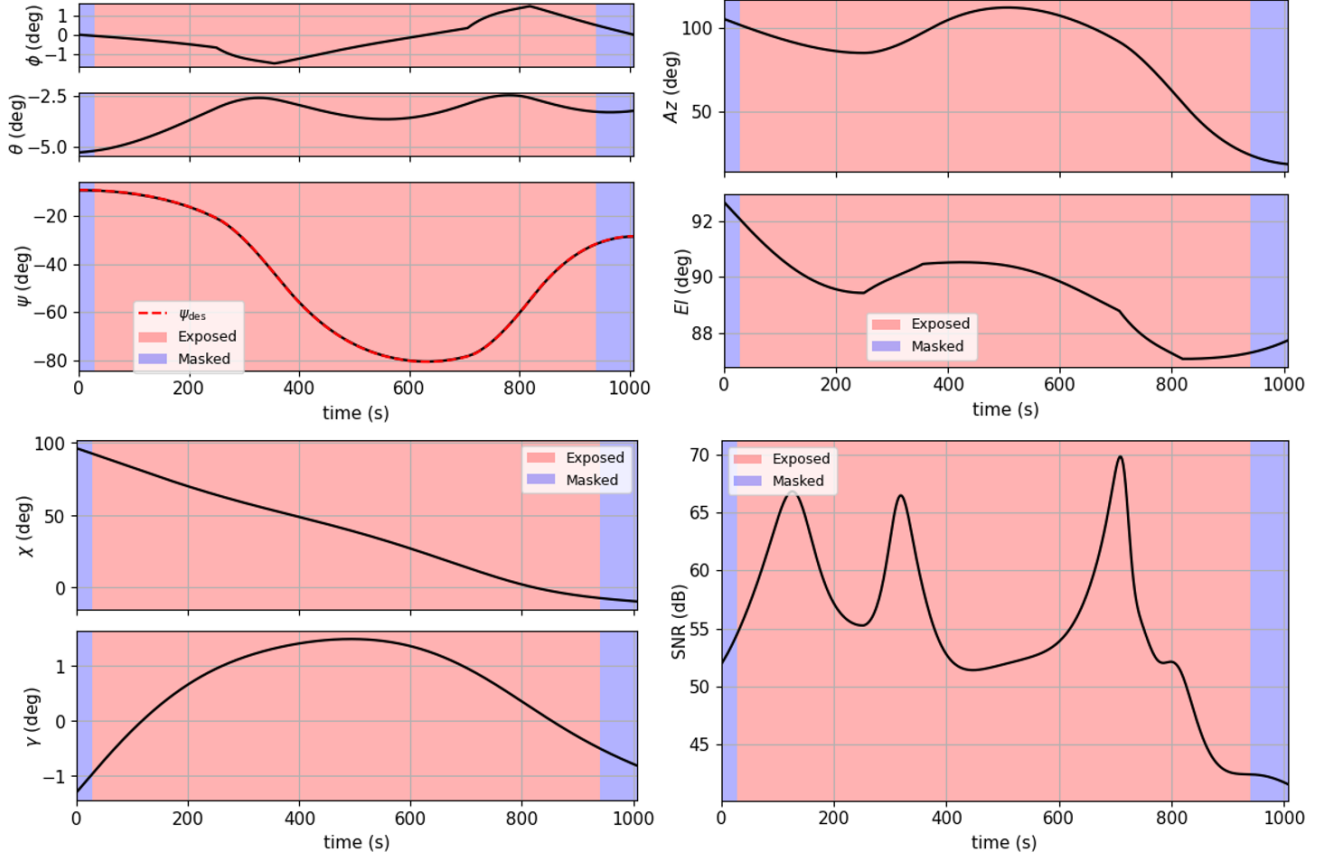
Case 3 is the joint path-attitude counterpart to Case 2, as shown in Fig. 19. The two cases use the same terrain masking weight, but Case 3 activates the SNR reduction term and allows attitude adjustment during exposed segments. In the time histories, ψ follows ψ_{des} mainly in terrain-masked segments, while the attitude and body-frame radar angles change during exposed segments to reduce the predicted SNR. Table 9 shows that Case 3 gives ER = 39.28% and $\tilde{S}_E = 47.47$ dB. Compared with Case 2, \tilde{S}_E is reduced by 5.49 dB.

The ER values and exposed intervals of Cases 2 and 3 are not exactly identical, so a small difference may exist when comparing their exposed-segment SNR values. Nevertheless, the two cases use the same terrain masking weight and exhibit similar ER values and major exposed intervals. Therefore, this comparison is meaningful for assessing the difference between path-only terrain masking and joint path-attitude optimization. Under these comparable exposure conditions, the reduction in \tilde{S}_E in Case 3 indicates that the proposed joint formulation lowers the surrogate-based detection metric not only by modifying the path, but also by exploiting attitude-dependent SNR patterns during exposure.

Table 9

Quantitative comparison of the simulation cases.

Case	ER (%)	\bar{S}_E (dB)	Trajectory length (m)	t_f (s)	Minimum WEZ margin (m)
Case 1	90.40	55.70	55182.00	1008.87	0.000012
Case 2	36.67	52.96	55111.43	1008.68	0.006222
Case 3	39.28	47.47	55133.81	1013.26	0.000401
Case 4	23.00	47.00	55053.20	1010.81	0.000001

**Figure 17:** Time histories of states, controls, and detection metrics for the baseline trajectory in Case 1 (without terrain masking or SNR reduction).

Case 4 is the joint path-attitude case with a larger terrain masking weight. As shown in Fig. 20 and Table 9, Case 4 reflects a stronger terrain masking preference and therefore exhibits a more occlusion-oriented flight tendency than Case 3. Quantitatively, Case 4 gives $ER = 23.00\%$ and $\bar{S}_E = 47.00$ dB. This indicates that increasing the terrain masking weight shifts the trajectory toward reduced exposure, while the SNR reduction term continues to guide attitude selection toward lower surrogate-predicted SNR during the remaining exposed segments.

WEZ constraint satisfaction is examined using Fig. 21 and Table 9. Fig. 21 plots the time histories of the horizontal separation $d_{h,j,k}$ and the corresponding maximum WEZ radius $\rho_{\max,j,k}$ for each SAM $j \in \mathcal{A} \triangleq \{1, 2\}$ in Cases 1–4. In all cases, $d_{h,j,k}$ remains greater than $\rho_{\max,j,k}$ throughout the simulated time horizon, indicating that the trajectories

stay outside the corresponding WEZ boundaries. This is also confirmed by the nonnegative minimum WEZ constraint residuals in Table 9, where values close to zero indicate active WEZ constraints within numerical tolerance.

To further interpret the attitude behavior during exposure, Cases 3 and 4 are examined in terms of the body-frame radar angles. In both cases, the UAV tends to maneuver toward operating regions around $Az \approx 110^\circ, 70^\circ, 30^\circ$ and $El \approx 87^\circ, 93^\circ$. These regions are consistent with the low-SNR patterns predicted by the SNR surrogate in Fig. 6, suggesting that the optimizer selects attitudes associated with lower surrogate-predicted SNR when exposure is unavoidable.

The resulting spatial behavior is further visualized in Figs. 22, 23, and 24. Fig. 24 provides a top-down view of the representative Case 4 trajectory relative to the terrain

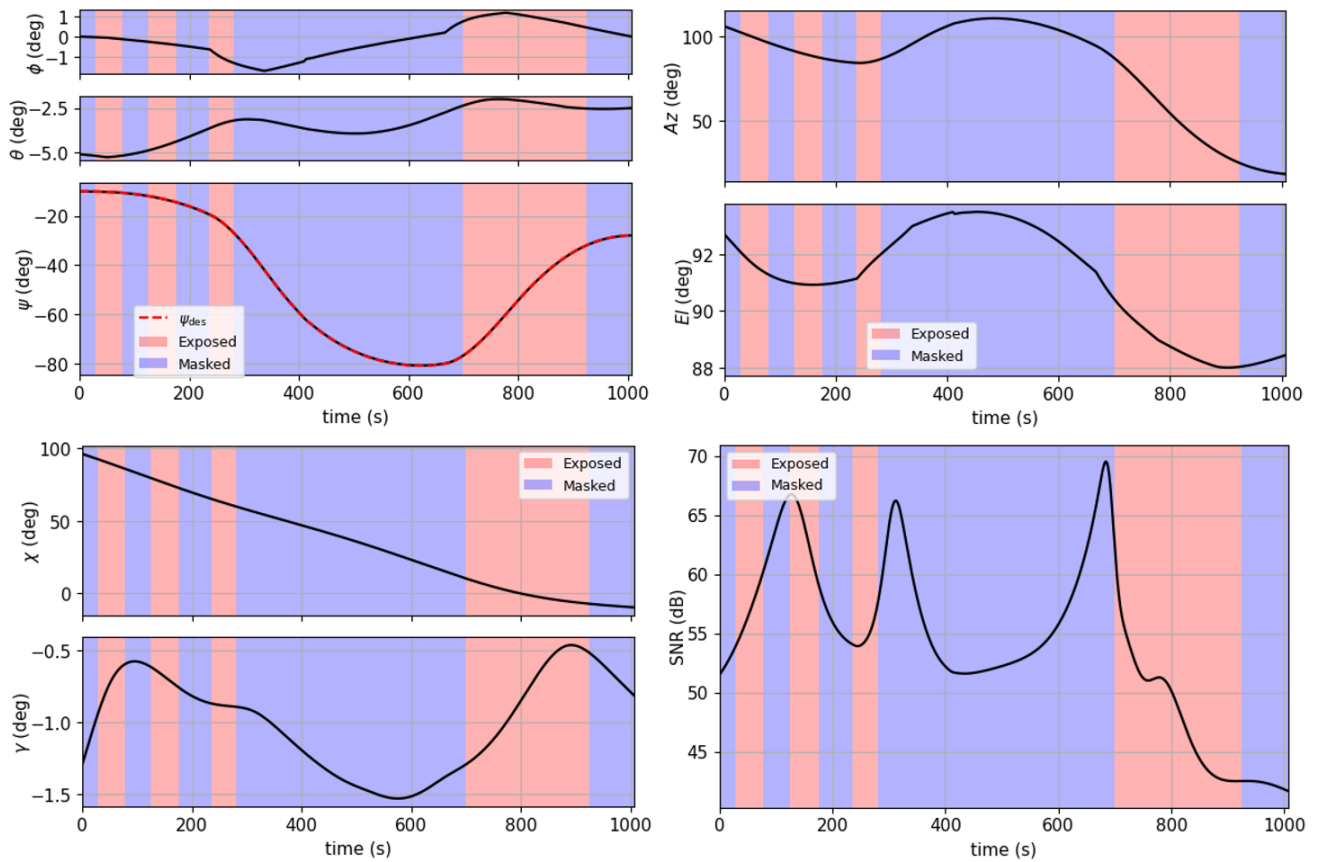


Figure 18: Time histories of states, controls, and detection metrics for the baseline trajectory in Case 2 (the path-only terrain masking baseline).

and WEZ regions. Fig. 22 and 23 present three-dimensional views of the baseline and representative low-observability cases over the mountainous terrain. In particular, Fig. 23 illustrates a trajectory that combines low-altitude terrain-following flight with attitude adjustments during exposed segments, consistent with the intended behavior of the proposed formulation.

6. Conclusion

This paper presented a coupled path-attitude optimization framework for low-observable UAV flight in complex terrain under radar and SAM threats. The formulation combines terrain masking, radar SNR reduction, and dynamic WEZ avoidance within a single trajectory design problem. By representing the radar SNR and the terrain exposure index with differentiable neural surrogate models, the framework embeds these nonlinear survivability metrics in an SCP-based optimization loop and obtains the required Jacobians efficiently through automatic differentiation.

In the representative simulation studies, the resulting trajectories satisfied the prescribed WEZ constraints, exploited terrain occlusion when available, and selected attitudes associated with lower surrogate-predicted SNR during exposed segments. The weight-sweep analysis further illustrated that the balance between terrain masking and SNR reduction

can be adjusted through the objective weights, providing a tunable trade-off across mission preferences.

This study should be interpreted as a formulation and proof-of-concept demonstration of coupled path-attitude optimization for low-observable flight. The surrogate fit results reported here are obtained on the constructed datasets, and the mission-level behavior is illustrated in a representative mountainous environment using selected trajectories from a weight-sweep study together with a baseline comparison case. Additional held-out generalization tests, exact-model post-validation, and closed-loop tracking evaluations remain important directions for future work. Future extensions will consider multistatic radar settings and additional observability channels, including EO/IR signatures.

Funding

This work was supported in part by Korea Research Institute for defense Technology planning and advancement (KRIT) grant funded by Defense Acquisition Program Administration (DAPA) (No. 20-105-E00-005(KRIT-CT-23-010), VTOL Technology Research Center for Defense Applications, 2025), and in part by KRIT Grant funded by DAPA (No. KRIT-CT-22-083, Cloud Computing-Based Learning/Inference Engine Technology). Equal support was received from both funding sources.

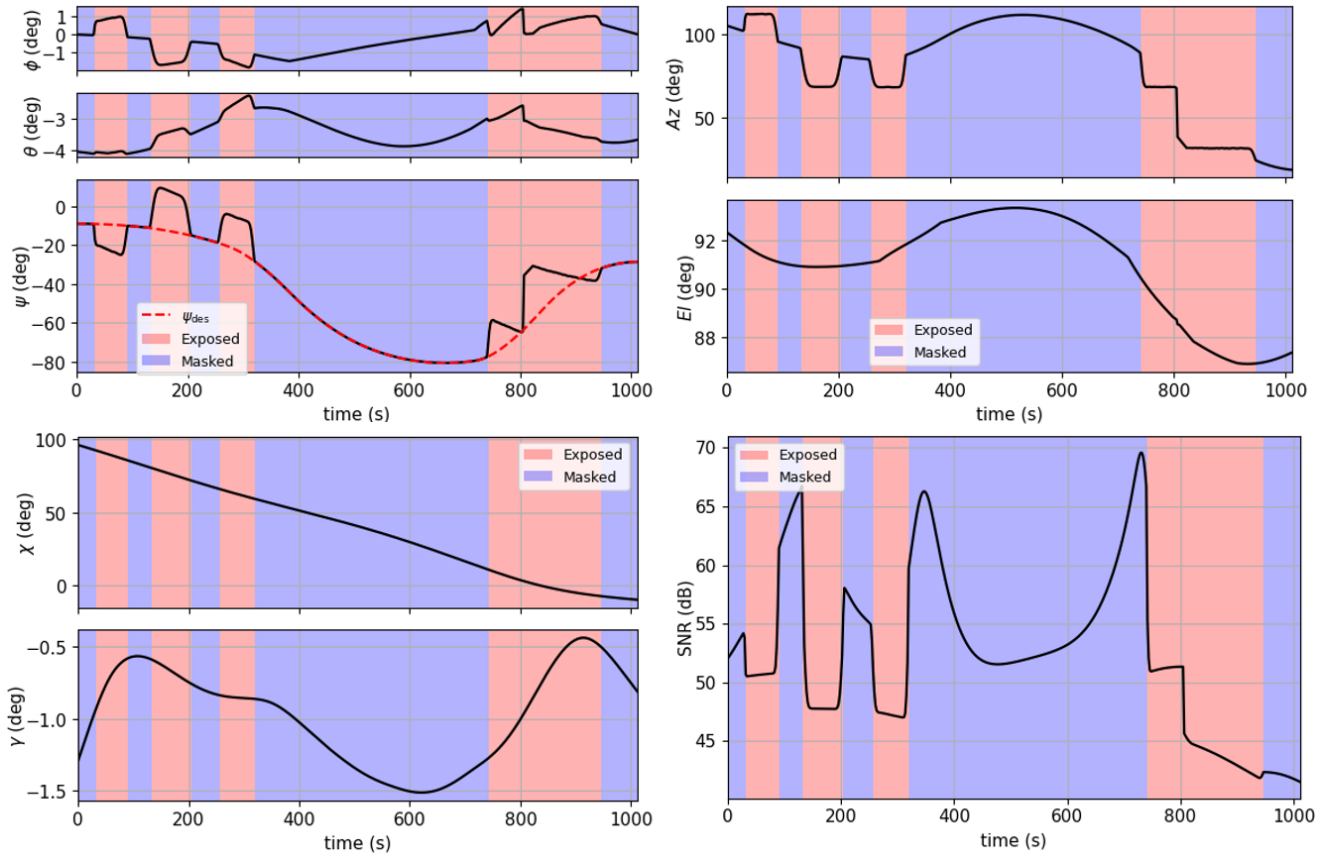


Figure 19: Time histories of states, controls, and detection metrics for representative Case 3 (smaller terrain masking weight).

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data and code that support the findings of this study are not publicly available because they include defense-related terrain, radar, and platform-signature data. Additional methodological details and selected derived data are available from the corresponding author upon reasonable request, subject to institutional and sponsor approval.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used Google Gemini to improve grammar, spelling, and phrasing, without changing technical contents. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

[1] T. Lapp, L. Singh, Model predictive control based trajectory optimization for nap-of-the-earth (NOE) flight including obstacle avoidance, in: Proceedings of the 2004 American Control Conference, Vol. 1, 2004, pp. 891–896 vol.1.

[2] A. Shirazi, J. Ceberio, J. A. Lozano, Spacecraft trajectory optimization: A review of models, objectives, approaches and solutions, Progress in Aerospace Sciences 102 (2018) 76–98.

[3] J. T. Betts, Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, Second Edition, 2nd Edition, Society for Industrial and Applied Mathematics, 2010.

[4] W. Meng, X. Zhang, L. Zhou, H. Guo, X. Hu, Advances in UAV path planning: A comprehensive review of methods, challenges, and future directions, Drones 9 (5) (2025) 376.

[5] M. Pelosi, C. Kopp, M. Brown, Range-limited UAV trajectory using terrain masking under radar detection risk, Applied Artificial Intelligence 26 (8) (2012) 743–759.

[6] Q. Xu, J. Ge, T. Yang, X. Sun, A trajectory design method for coupling aircraft radar cross-section characteristics, Aerospace Science and Technology 98 (2020) 105653.

[7] J. Hu, L. Fan, Y. Lei, Z. Xu, W. Fu, G. Xu, Reinforcement learning-based low-altitude path planning for uas swarm in diverse threat environments, Drones 7 (9) (2023).

[8] P. M. Dillon, M. D. Zollars, I. E. Weintraub, A. Von Moll, Optimal trajectories for aircraft avoidance of multiple weapon engagement zones, Journal of Aerospace Information Systems 20 (8) (2023) 520–525.

[9] C. Sánchez-Sánchez, D. Izzo, Real-time optimal control via deep neural networks: Study on landing problems, Journal of Guidance, Control, and Dynamics 41 (5) (2018) 1122–1135.

[10] W. Li, W. Li, L. Cheng, S. Gong, Trajectory optimization with complex obstacle avoidance constraints via homotopy network sequential convex programming, Aerospace 9 (11) (2022) 720.

[11] T. A. Almseidein, A. Alzidaneen, Optimizing UAV trajectories with multi-layer artificial neural networks, Informatica 49 (2) (2025).

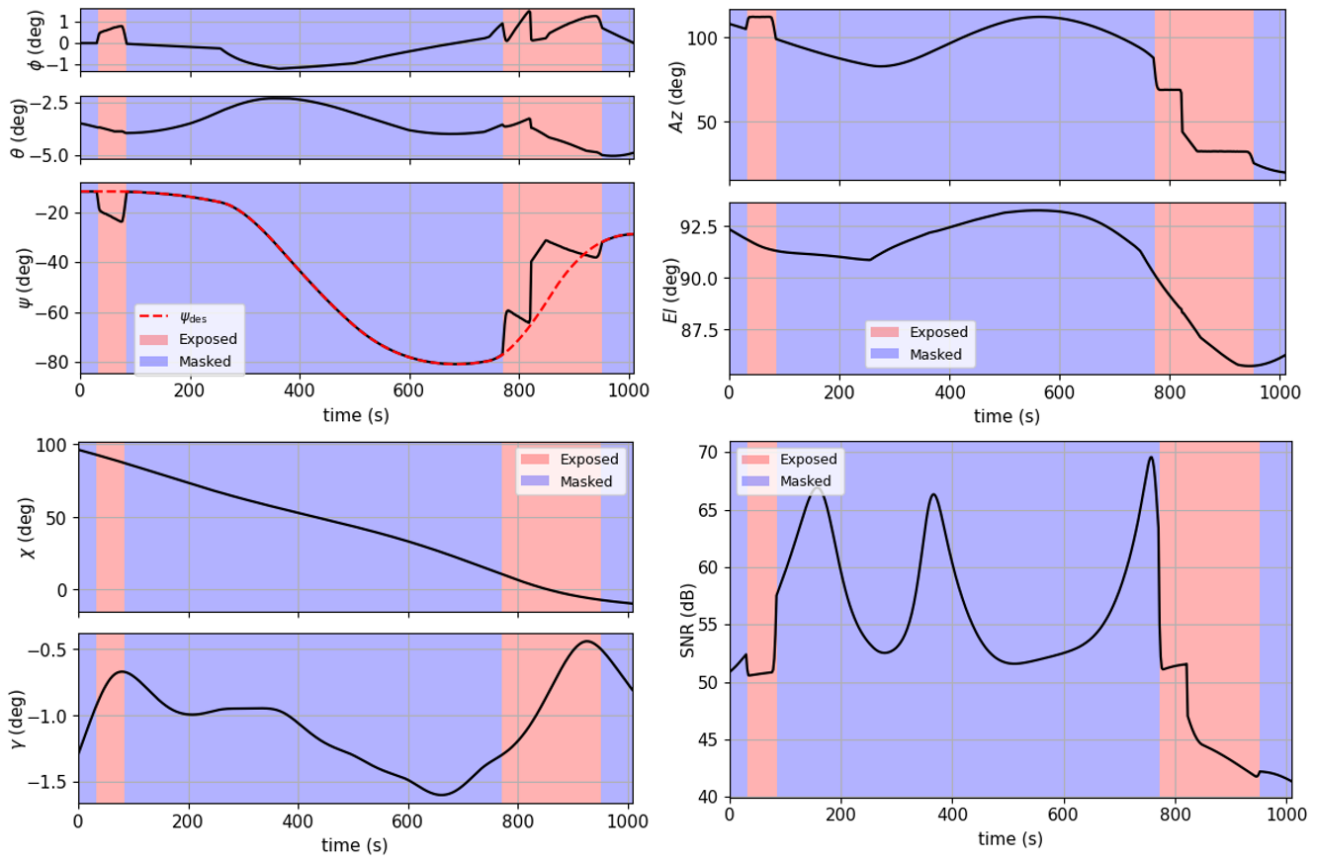


Figure 20: Time histories of states, controls, and detection metrics for representative Case 4 (larger terrain masking weight).




Figure 21: SAM-relative horizontal distance and maximum WEZ radius over time for Cases 1–4.

- [12] M. Eskandari, A. V. Savkin, M. Deghat, Kinodynamic model-based UAV trajectory optimization for wireless communication support of internet of vehicles in smart cities, *Drones* 8 (10) (2024) 574.
- [13] Y. Mao, M. Szmuk, B. Açıkmeşe, Successive convexification of non-convex optimal control problems and its convergence properties, in: 2016 IEEE 55th Conference on Decision and Control (CDC), 2016, pp. 3636–3641.
- [14] Z. Wang, M. J. Grant, Constrained trajectory optimization for planetary entry via sequential convex programming, *Journal of Guidance, Control, and Dynamics* 40 (10) (2017) 2603–2615.
- [15] A. Hakobyan, I. Yang, Wasserstein distributionally robust motion control for collision avoidance using conditional value-at-risk, *IEEE Transactions on Robotics* 38 (2) (2021) 939–957.
- [16] S. Xu, G. Li, J. Liu, Obstacle avoidance for manipulator with arbitrary arm shape using signed distance function, in: 2018 IEEE International Conference on Robotics and Biomimetics (ROBIO), IEEE, 2018, pp. 343–348.
- [17] U.S. Army, Operator’s Manual for UH-60A, UH-60L, and EH-60A Helicopters, Department of the Army, TM 1-1520-237-10 (2009).

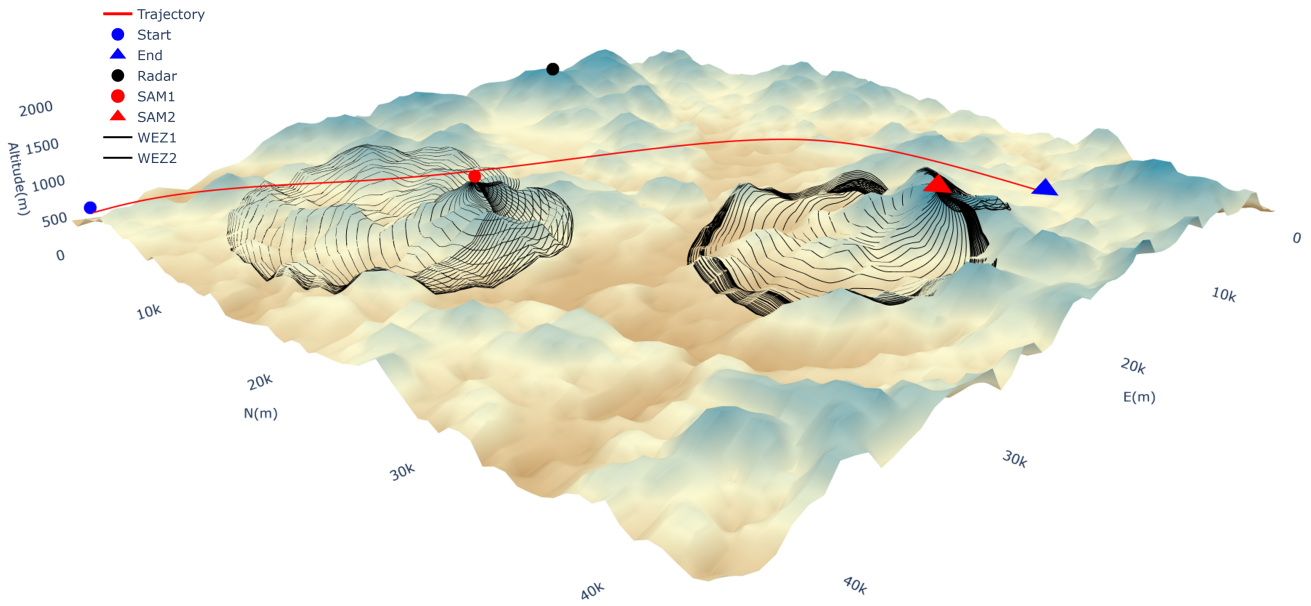


Figure 22: Three-dimensional view of the baseline Case 1 trajectory in the mountainous terrain.

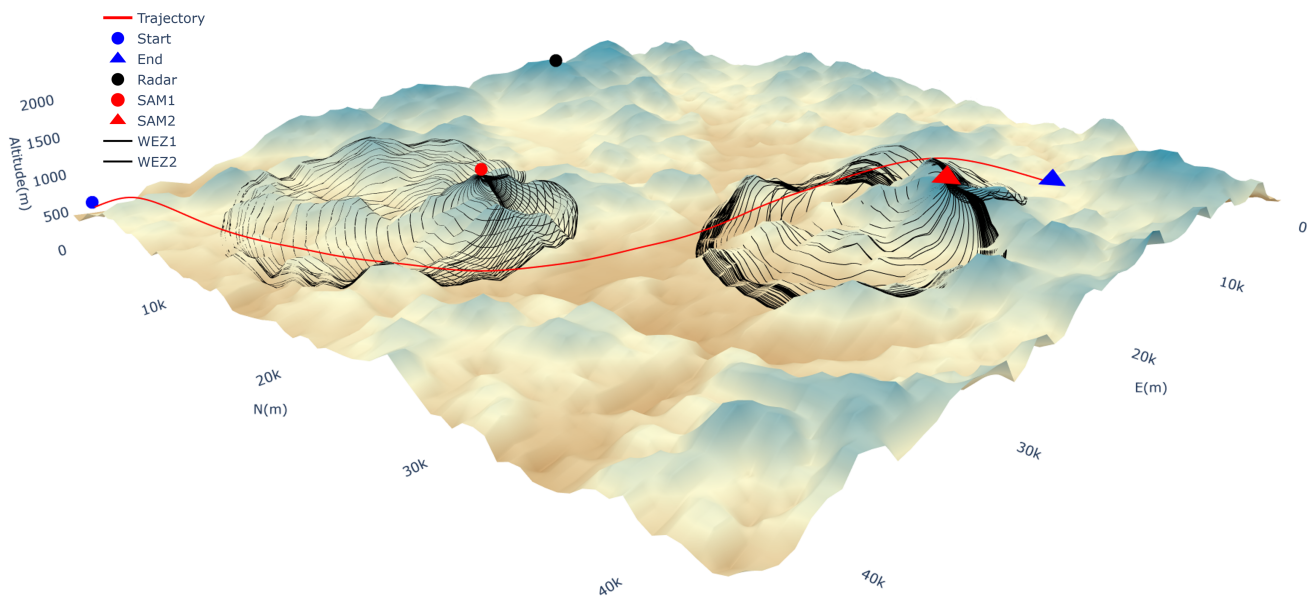


Figure 23: Three-dimensional view of the representative Case 4 trajectory in the mountainous terrain.

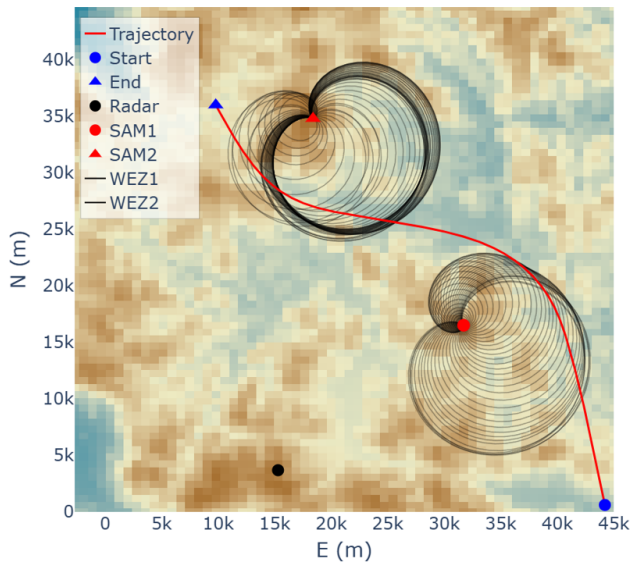


Figure 24: Top-down view of the representative Case 4 trajectory relative to terrain and WEZ regions.

A. Nomenclature

Symbol	Description
X	UAV state vector
U	UAV control input vector
$p = [p_N, p_E, p_D]^T$	UAV position in the NED frame
$v = [v_N, v_E, v_D]^T$	UAV velocity in the NED frame
ϕ, θ, ψ	Roll, pitch, and yaw angles
$T = [T_N, T_E, T_D]^T$	Thrust vector in the NED frame
ω	Yaw rate
m	UAV mass
g	Gravitational acceleration vector
$D(v)$	Aerodynamic drag force
ρ	Air density
C_D	Drag coefficient
S_{ref}	Reference area for aerodynamic drag calculation
t_f	Terminal time
τ	Normalized time
A_k, B_k, c_k	Linearized continuous-time dynamics matrices
Φ_k, Γ_k	Discrete-time state transition and input matrices
Az, El	Body-frame radar azimuth and elevation angles
χ, γ	Radar-centric LOS azimuth and elevation angles
r	Radar-UAV LOS vector
d	Radar-UAV slant range
σ_{dB}	RCS in dB scale
S, \hat{S}	Ground-truth radar SNR and surrogate-predicted radar SNR
E, \hat{E}	Ground-truth and predicted terrain exposure index
\hat{E}^{rob}	Robust terrain exposure prediction with CVaR margin
f_S	SNR neural surrogate model
f_E	Terrain exposure neural surrogate model
x_S, x_E	Input feature vectors for the SNR and exposure surrogates
γ_{sky}	Skyline elevation angle
$\Delta\gamma$	Elevation-angle offset relative to the skyline
m_{FN}	CVaR-based false-negative safety margin
α	CVaR confidence level
σ_k^E, σ_k^S	Switching coefficients for terrain-masked and exposed-regimes
\mathcal{A}	Set of SAM systems
$p_{\text{SAM},j}$	Position of SAM site j
$p_{h,k}$	UAV horizontal position at time step k
$v_{h,k}$	UAV horizontal velocity at time step k
$d_{h,j,k}$	Horizontal distance between the UAV and SAM j
θ_h	SAM launcher pointing direction
χ_h	Horizontal LOS bearing from the SAM to the UAV
$\xi_{j,k}$	Relative angle between UAV horizontal velocity and SAM-UAV LOS direction
ρ_{max}	Maximum cardioid-shaped WEZ radius
R_{max}	Maximum horizontal engagement range of the SAM
\mathcal{O}	Set of terrain facets
$p_{\text{facet},i}$	Centroid of terrain facet i
n_i	Outward unit normal vector of terrain facet i
z_{safe}	Minimum terrain clearance distance
$\psi_{\text{des},k}$	Desired heading angle aligned with ground-projected velocity
$\Delta_{t_f}, \Delta_p, \Delta_v, \Delta_\psi$	Trust-region radii for terminal time, position, velocity, and yaw

B. SCP-Based Computational Procedure

To clarify the implementation sequence, Algorithm 1 summarizes the SCP-based computational procedure used in the proposed coupled path-attitude optimization framework. The algorithm follows the notation and equations introduced in Sections 2–4.

Algorithm 1 SCP-Based Coupled Path–Attitude Optimization Procedure

Require: Terrain facets \mathcal{O} , radar position p_{radar} , SAM positions $\{p_{\text{SAM},j}\}_{j \in \mathcal{A}}$, mission endpoints p_{init} and p_{fin} , trained surrogates f_S and f_E , CVaR margin m_{FN} , and parameters in Tables 1 and 7

Ensure: Optimized trajectory (X^*, U^*, t_f^*)

- 1: Define $\hat{E}^{\text{rob}} = \hat{E} + m_{\text{FN}}$ using (26).
 - 2: Load or construct a warm-start trajectory (X^0, U^0, t_f^0) satisfying $p_0 = p_{\text{init}}$ and $p_N = p_{\text{fin}}$.
 - 3: Initialize trust-region radii Δ_{t_f} , Δ_p , Δ_v , and Δ_ψ .
 - 4: **for** $n = 0, 1, \dots$ **do**
 - 5: Set the reference trajectory $(\bar{X}, \bar{U}, \bar{t}_f) \leftarrow (X^n, U^n, t_f^n)$.
 - 6: **for** $k = 0, \dots, N - 1$ **do**
 - 7: Compute $D(\bar{v}_k)$ and $J_{v,k}^D$ using (3)–(4).
 - 8: Construct (8) and compute $\Phi_k, \Gamma_{u,k}, \Gamma_{c,k}, \Gamma_{s,k}$ using (15).
 - 9: Reconstruct (ϕ_k, θ_k) using (16).
 - 10: Evaluate (Az_k, El_k) , (χ_k, γ_k) , and d_k from Secs. 2.3–2.4.
 - 11: Evaluate $\hat{S}_k = f_S(Az_k, El_k, d_k)$ and $\hat{E}_k^{\text{rob}} = f_E(\chi_k, \gamma_k, d_k) + m_{\text{FN}}$.
 - 12: Compute σ_k^E and σ_k^S using (39).
 - 13: Compute $\rho_{\max,j,k}$ using (27)–(28), for all $j \in \mathcal{A}$.
 - 14: Linearize (Az_k, El_k) , (χ_k, γ_k) , \hat{S}_k , \hat{E}_k , and $\rho_{\max,j,k}$ using (34)–(38).
 - 15: Select adjacent terrain facets \mathcal{O}_k and construct the terrain clearance constraint (33).
 - 16: **end for**
 - 17: Form the convex SCP subproblem (41) using trust-region bounds (40).
 - 18: Solve (41) using GUROBI and obtain $(X^{\text{cand}}, U^{\text{cand}}, t_f^{\text{cand}})$.
 - 19: Evaluate the predicted reduction $\Delta \mathcal{J}_{\text{pred}}$ and actual nonlinear reduction $\Delta \mathcal{J}_{\text{act}}$.
 - 20: Compute the agreement ratio $\rho_{\text{SCP}} = \Delta \mathcal{J}_{\text{act}} / \Delta \mathcal{J}_{\text{pred}}$.
 - 21: Compute the linearization error δ_{lin} .
 - 22: **if** $\rho_{\text{SCP}} \geq \rho_{\min}$ and $\delta_{\text{lin}} \leq \delta_{\max}$ **then**
 - 23: Accept the candidate and set $(X^{n+1}, U^{n+1}, t_f^{n+1}) \leftarrow (X^{\text{cand}}, U^{\text{cand}}, t_f^{\text{cand}})$.
 - 24: Adapt Δ_{t_f} , Δ_p , Δ_v , and Δ_ψ .
 - 25: **else**
 - 26: Reject the candidate and set $(X^{n+1}, U^{n+1}, t_f^{n+1}) \leftarrow (X^n, U^n, t_f^n)$.
 - 27: reduce Δ_{t_f} , Δ_p , Δ_v , and Δ_ψ .
 - 28: **end if**
 - 29: Compute the trajectory update error $\epsilon_X = \|X^{n+1} - X^n\|_F$ and objective variation $\epsilon_J = |\Delta \mathcal{J}_{\text{act}}|$.
 - 30: **if** $\epsilon_X \leq \epsilon_X^{\text{tol}}$ and $\epsilon_J \leq \epsilon_J^{\text{tol}}$ **then**
 - 31: **break**
 - 32: **end if**
 - 33: **end for**
 - 34: **return** (X^*, U^*, t_f^*) .
-