

# HALO: Hazard-Aware Landing Optimization for Autonomous Systems

Christopher R. Hayner<sup>\*</sup>, Samuel C. Buckner<sup>\*†</sup>, Daniel Broyles<sup>‡</sup>,  
 Evelyn Madewell, Karen Leung, Behçet Açıkkmeş

**Abstract**—With autonomous aerial vehicles enacting safety-critical missions, such as the Mars Science Laboratory Curiosity rover’s landing on Mars, the tasks of automatically identifying and reasoning about potentially hazardous landing sites is paramount. This paper presents a coupled perception-planning solution which addresses the hazard detection, optimal landing trajectory generation, and contingency planning challenges encountered when landing in uncertain environments. Specifically, we develop and combine two novel algorithms, Hazard-Aware Landing Site Selection (HALSS) and Adaptive Deferred-Decision Trajectory Optimization (ADDDTO), to address the perception and planning challenges, respectively. The HALSS framework processes point cloud information to identify feasible safe landing zones, while ADDDTO is a multi-target contingency planner that adaptively replans as new perception information is received. We demonstrate the efficacy of our approach using a simulated Martian environment and show that our coupled perception-planning method achieves greater landing success whilst being more fuel efficient compared to a non-adaptive DDTO approach.

## I. INTRODUCTION

Autonomous robotic systems are poised to carry out a number of safety-critical missions, such as planetary exploration and disaster response, that otherwise would be dangerous, expensive, or impractical for humans to perform. In the case of autonomous aerial robots (e.g., unoccupied aerial vehicles), where landing safely is critical to mission success, the robots must reason about terrain conditions to identify and avoid potentially hazardous landing areas which threaten the mission goals.

In this work, we propose a perception-aware trajectory planning algorithm that enables an autonomous aerial robot to land safely in unknown or uncertain environments. Specifically, our algorithm combines a LiDAR-based hazard-aware landing site selector with an adaptive multi-target trajectory optimizer to generate fuel-efficient trajectories that are robust to changes in landing sites priority order as new observations are received.

**Related work.** Hazard detection is a critical element of NASA’s efforts to develop precision landing technologies for planetary exploration missions [1]. The NASA Safe

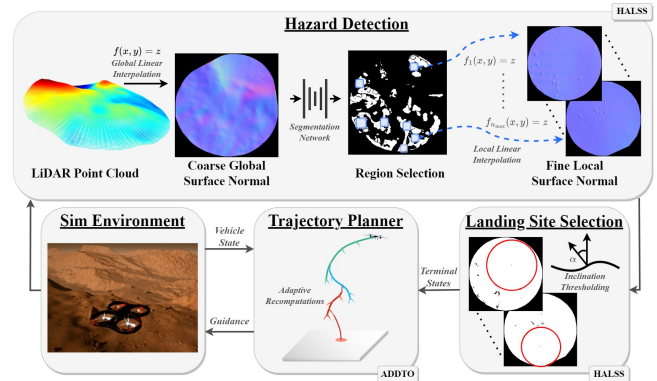


Fig. 1: System Overview: HALSS identifies viable landing sites used in the multi-target trajectory optimization generated by ADDDTO. The resulting guidance is used to command a (simulated) quadrotor, and landing site viability is continually evaluated as new measurements are processed.

and Precise Landing - Integrated Capabilities Evolution (SPLICE) project specifically outlines next-generation hazard detection capabilities, which primarily operate on lidar-derived digital elevation models (DEMs) to measure the slope and roughness of potential landing sites [2]. These DEM-based geometric approaches to hazard classification include probabilistic modeling [3]–[5] and fuzzy-reasoning-based safety maps [6]. However recent work has achieved performance improvements over classical methods via the use of deep learning, including landing site hazard identification from semantic segmentation networks [7], [8] and Bayesian deep-learned segmentation and prediction variance for uncertainty-aware safe landing site identification [9]. Our proposed method leverages the benefits of Bayesian segmentation with the theoretical guarantees of a direct geometric representation of hazardous terrain.

An important consideration in perception-aware trajectory planning is the management of contingency plans when a previously-viable target state is found to be infeasible or impractical due to map and localization uncertainty, fine-grain objects not initially detected by LiDAR or other on-board sensors, and uncertainty in the environment. Model predictive control (MPC) approaches have demonstrated robust performance for trajectory guidance in dynamic and complex landing scenarios [10], however MPC does not consider multiple contingency options. Deferred-decision trajectory optimization (DDTO) [11] offers a framework to compute trajectories to multiple selected targets while maximizing the deferral time before needing to commit to a target for a given

<sup>\*</sup>Both authors contributed equally.

<sup>†</sup> Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

<sup>‡</sup> The views expressed in this article are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

The authors would like to acknowledge Annika Singh and Michael Hayner.

Dept. of Aeronautics and Astronautics, University of Washington, USA. Emails: {haynec, sbuckne1}@uw.edu

suboptimality tolerance. DDTO is notably distinguished from stochastic methods to handle uncertainty in that it is entirely deterministic and requires no probabilistic knowledge of the environment. The trajectory generation method presented in this paper (ADAPTIVE-DDTO) combines the advantages of both DDTO and MPC.

**Contributions.** Contributions of this work include novel frameworks for determining safe landing sites in uncertain and hazardous areas, known as Hazard-Aware Landing Site Selection (HALSS), and for adaptive multi-target trajectory optimization, referred to as (ADAPTIVE-DDTO).

**Organization.** Sections II and V-A discuss the proposed system architecture and simulation environment used for this research, respectively. The HALSS framework is described in Section III, and Section IV covers the ADAPTIVE-DDTO algorithm. Results from an integrated simulation and comparative study are presented in Section V, followed by conclusions in Section VI.

## II. ARCHITECTURE

Fig. 1 depicts our proposed solution at a high level. The HALSS algorithm performs the hazard detection from LiDAR measurements of the terrain and determines safe landing sites represented by circular regions. These landing sites are passed as terminal states to ADAPTIVE-DDTO, which removes landing sites as they are found to be unsafe or are no longer reachable as the vehicle passes by its respective branch point in the trajectory. ADAPTIVE-DDTO then solves the landing optimization problem, and the resulting multi-target trajectory is executed by a low-level controller. The landing sites are returned to HALSS to have their radii updated if any new unsafe terrain are identified within. This loop is repeated until the autonomous system reaches a cutoff altitude, at which point it locks to the best landing site.

### III. HAZARD-AWARE LANDING SITE SELECTION

In this section, we introduce the HALSS framework that selects safe landing sites from point cloud data. Notably, HALSS performs hazard detection at multiple scales which allows for efficient computation without sacrificing hazard detection performance. We define the global scale as encompassing the entirety of the observed map and local scale as smaller subsets of the observed map. HALSS first identifies larger-scale hazards by performing coarse hazard detection on the global scale through the process outlined in III-A. From the resulting global safety map, it selects regions to identify potential landing sites using the process outlined in III-C. For each region identified by the coarse hazard detection, fine hazard detection is performed to identify smaller-scale hazards using the process outlined in III-B. From each of the identified local safety maps, a landing site is selected using the process outlined in III-C.

#### A. Coarse Hazard Detection

Coarse hazard detection aims to identify landing hazards, such as large slopes and rocks, in order to generate a global safety map. This task is accomplished via a

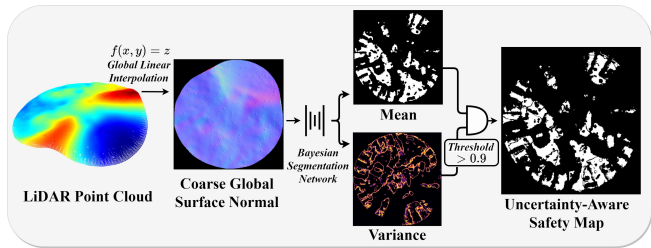


Fig. 2: Coarse Hazard Detection Process: a LiDAR-derived surface normal image is fed into a Bayesian SegNet which classifies flat and sloped/rocky regions; the resulting output is an uncertainty-aware safety map with white pixels corresponding to safe/flat areas.

Bayesian segmentation network for classifying terrain as safe/unsafe. This section discusses the coarse hazard detection pipeline, as shown in Fig. 2.

*Global Surface Normal Interpolation:* The surface normal representation of the terrain is used as the input to the segmentation network due to its ability to express terrain slope and roughness features. It consists of three-channel data which encodes the unit surface normal vector for each pixel and can be derived from LiDAR point cloud data (as shown in Fig. 2). To create a global surface normal image from a point cloud, we first cull points to decrease the interpolation time using a roughness-preserving downsampling algorithm. This algorithm constructs a two-dimensional uniform grid where, for each cell, the points with the largest and the smallest Z-value are chosen, and the remaining points are culled. Choosing the largest and smallest Z-value points ensures that the roughness is preserved in the culled point cloud, which is the most conservative case and results in the least safe terrain. A function is then linearly interpolated from the culled point cloud and evaluated at uniform grid spacing.

*Network Architecture:* We use a Bayesian segmentation network based on the Bayesian Seg-Net architecture [12], which is well suited to learn higher dimensional features (such as larger slopes and rocks). Monte Carlo (MC) dropout layers are included after each series of convolution operations to account for uncertainty in the predictions. To produce a binary safety map, we apply a sigmoid transform to the output of the network. During inference, the global surface normal is propagated through the network, resulting in 30 stochastic samples. The mean and variance of these outputs model the predicted safety and approximated Bayesian uncertainty of that prediction, respectively, similar to the process outlined in [9]. We then normalize and threshold the variance at a conservative value of 0.9 which was experimentally chosen and apply the logical AND operation to the variance and mean, resulting in the uncertainty-aware safety map.

#### B. Fine Hazard Detection

Potential landing sites identified from the global safety map are further analyzed to identify small hazards and create detailed, high-resolution safety maps for each landing region. This section explains the fine hazard detection process, as depicted in Fig. 3.

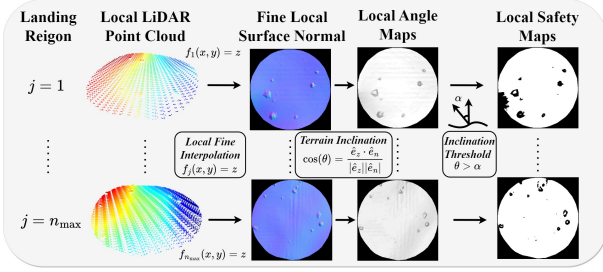


Fig. 3: Fine Hazard Detection Process: local safety maps are obtained via angle thresholding high-resolution local surface normals for each landing region.

*Localized Surface Normal Interpolation:* In order to ensure that smaller localized details are captured with respect to each region, a subset of the original un-culled point cloud within the radius (plus a buffer) of the center coordinate of the identified regions is taken, resulting in a set of localized point clouds, as seen in Fig. 3. To prevent the computation time from growing unbounded as new points are added to the global point cloud, each of these point clouds is then independently downsampled using the same roughness-preserving downsampling technique introduced in III-A, but with finer grid spacing. Similar to the global surface normal, the localized surface normals are created by linearly interpolating a function to each culled local point cloud, sampling at even grid spacing, and finally calculating the surface normals.

*Inclination Thresholding:* To identify smaller and more low-dimensional hazards, a pure geometric representation of the slope of the terrain is well suited as it offers explainable representations of smaller hazards. The local angle map is obtained by taking the dot product between the unit surface normal vector,  $\hat{e}_n \in \mathbb{R}^3$ , from each local point cloud with the vertical unit vector,  $\hat{e}_z \in \mathbb{R}^3$ , at each point in the local surface normal map. The local angle map,  $\theta \in \mathbb{R}$ , is then thresholded against the maximum allowable inclination,  $\alpha \in \mathbb{R}$ , such that  $\theta > \alpha$  implies an unsafe region, which results in the local safety maps.

### C. Region and Landing Site Selection

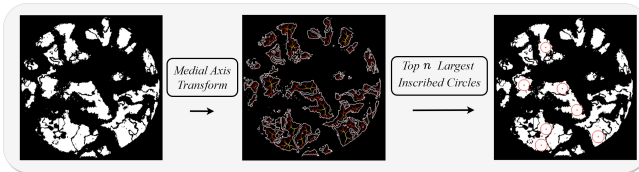


Fig. 4: Example of safe region (red circles) identification via medial axis skeletonization of the binary safety map.

Potential safe regions and landing sites are identified by considering the largest inscribed circle contained in the safe region, which is easily obtained from the medial axis transform (MAT) [13]. The medial axis is a continuous version of the Voronoi diagram, and in the 2-D domain can be thought of as the set of centers of maximally inscribed discs with respect to some boundary. Applying the MAT to a binary image results in a skeleton, which is an array of values representing the distance to a boundary in the original object; the boundary in this

particular case is the set of black pixels. Fig. 4 shows an example of medial axis skeletonization of a safety map. To find a set of candidate landing sites, the global safety map skeleton is used to quickly and efficiently find the top  $n$  (number of desired landing targets) largest circular safe zones. For fine hazard detection, this process is repeated for each local safety map with  $n = 1$ .

*Scoring Parameters:* Once obtained, each landing site is scored using the following metrics: the point cloud density within the landing site (*density*), the Euclidean distance between the landing site and vehicle in the NED frame (*proximity*), the proximity to other landing sites (*clustering*), the 99<sup>th</sup> percentile of the uncertainty map within landing site (*uncertainty*), and the landing site bounding circle radius (*size*).

## IV. ADAPTIVE DEFERRED-DECISION TRAJECTORY OPTIMIZATION

In the context of hazard-awareness, acquired targets can lose viability at a future time due to navigation errors and environmental uncertainty, thus reducing the set of available targets. In this section, a novel algorithm is introduced to maintain a bounded number of targets in the reachable set of the vehicle for path planning at all times. This approach starts with (1) the baseline trajectory optimization problem formulation, and builds on it with (2) a single-shot multi-targeted path planning approach (DDTO), which is enhanced to an adaptively-recomputed multi-targeted path planning approach (ADAPTIVE-DDTO).

### A. Baseline Problem Formulation

To model the vehicle dynamics (in the context of a multirotor), a three degree-of-freedom (3-DOF) state-space model is considered of the form:

$$\dot{x}(t) = Ax(t) + Bu(t) + p \quad (1)$$

where  $r(t) \in \mathbb{R}^3$  is the vehicle's position,  $\dot{r}(t) \in \mathbb{R}^3$  is the vehicle's velocity, and  $x(t) = [r(t)^\top, \dot{r}(t)^\top]^\top \in \mathbb{R}^6$  is the vehicle's state. Furthermore, the control signal  $u(t) \in \mathbb{R}^3$  is the commanded thrust of the vehicle (henceforth denoted as  $T(t)$ ). Consequently, the vehicle's simplified dynamics can be described as a double integrator system with an affine gravity term, such that:

$$A = \begin{bmatrix} 0 & I \\ 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ \frac{I}{m} \end{bmatrix}, \quad p = \begin{bmatrix} 0 \\ -g \end{bmatrix} \quad (2)$$

where  $g \in \mathbb{R}^3$  is the acceleration due to gravity,  $m \in \mathbb{R}$  is the mass of the vehicle, and  $I \in \mathbb{R}^{3 \times 3}$  is the identity matrix. We consider a surface-fixed inertial coordinate frame consisting of basis vectors  $e_j \in \mathbb{R}^3$ , for  $j = 1, 2, 3$ , which correspond to downrange, crossrange and altitude, respectively. The state-space model (2) assumes that gravity is constant and neglects aerodynamic drag.

The objective of the single-shot trajectory optimization problem (BASEPROBLEM) is to guide the target from its current airborne state,  $x_0$ , to a safe target zone on the ground,  $x_f$ , subject to the dynamics in (1), as well as constraints on the state  $x(t)$  and control  $u(t)$ , while minimizing control effort (cumulative thrust):

BASEPROBLEM (3-DOF Non-Convex)	(3)
$\underset{T(\cdot)}{\text{minimize}} \quad \int_0^{t_f} \ T(\tau)\ _2 d\tau$	(3a)
$\text{subject to} \quad \dot{x}(t) = Ax(t) + Bu(t) + p,$	(3b)
$0 < T_{\min} \leq \ T(t)\ _2 \leq T_{\max},$	(3c)
$\ T(t)\ _2 \cos(\gamma_{\max}) \leq e_3^\top T(t),$	(3d)
$\ \dot{r}_{1:2}(t)\ _2 \leq v_{\max}^{\text{lat}},$	(3e)
$\text{abs}(\dot{r}_3(t)) \leq v_{\max}^{\text{vert}},$	(3f)
$x(0) = x_0, x(t_f) = x_f$	(3g)

In problem (3), cumulative commanded thrust is minimized for in the cost function (3a), with constraints for dynamics (3b), lower ( $T_{\min}$ ) and upper ( $T_{\max}$ ) bound constraints on thrust magnitude (3c), a tilt angle bound  $\gamma_{\max}$  to the vertical  $e_3$  (3d), a lateral velocity bound  $v_{\max}^{\text{lat}}$  (3e), a vertical velocity bound  $v_{\max}^{\text{vert}}$  (3f), and initial and terminal boundary conditions (3g), across a fixed time-of-flight denoted  $t_f \in \mathbb{R}$ .

The lower bound in (3c) causes the constraint, and therefore the problem, to become non-convex. The lossless convexification (LCvx) approach (see [14], [15]) provides the capability to convexify these constraints in a way that is both lossless (preserves global optimality) and provides feasibility guarantees. This can be done by introducing the slack variable  $\Gamma(t) \in \mathbb{R}$ , which substitutes  $\|T(t)\|_2$  in the cost (3a) and constraints (3c)–(3d) and provides an additional relaxation constraint  $\|T(t)\|_2 \leq \Gamma(t)$ . This augmentation to BASEPROBLEM will be used throughout the remainder of this paper and referred to as LCvxPROBLEM. The 3-DOF solution that LCvxPROBLEM provides can be considered a proxy to a full 6-DOF solution (with yaw unconstrained) due to the differential flatness property of multirotor vehicles [16], [17].

### B. Adaptive Deferred-Decision Trajectory Optimization

Deferred-decision trajectory optimization (DDTO) [11] offers an enhancement to LCvxPROBLEM to enable multi-target path planning while keeping multiple target states in the reachable set of the vehicle and maximizing the deferral time to each of them given a set of suboptimality tolerances (relative to the optimal single-shot LCvxPROBLEM solution to the same target) and a target rejection sequence (see Fig. 5). The ADAPTIVE-DDTO algorithm builds on this idea to provide a fault-tolerant trajectory optimization framework to enable the vehicle to maintain at least  $n_{\min} \in \mathbb{N}$  targets and up to  $n_{\max} \in \mathbb{N}$  targets in consideration at all times (see Fig. 6 and Algo. 1). The following design can be considered similar to a shrinking-horizon MPC problem (e.g., [10]) which is conditionally executed after non-uniform time intervals. The following paragraphs will describe the prerequisite knowledge behind Algorithm 1.<sup>1</sup>

<sup>1</sup>Note that in Algorithm 1, when calling the underlying DDTO algorithm described in [11], LCvxPROBLEM is used as the OPT-PROBLEM, with the corresponding constraint set used to solve each SUBPROBLEM.

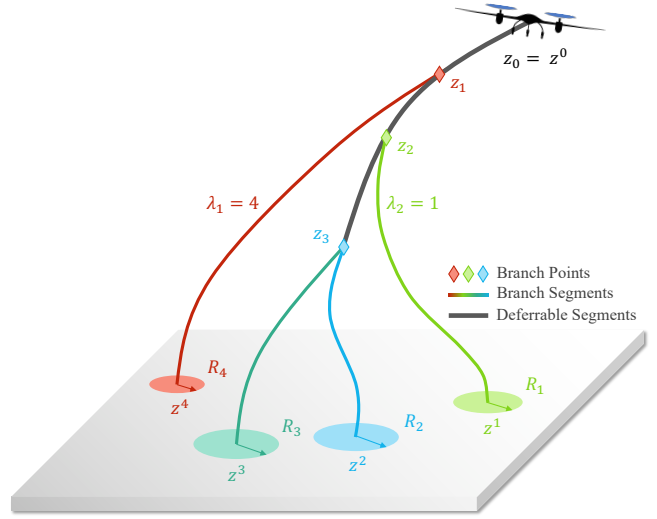


Fig. 5: Illustration of a DDTO solution to the multirotor landing problem with  $n = 4$  targets.

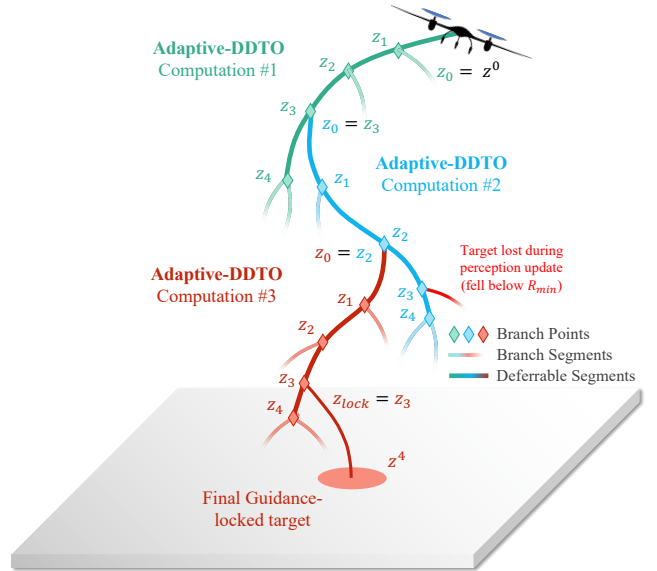


Fig. 6: Illustration of an ADAPTIVE-DDTO solution to the multirotor landing problem with  $n_{\min} = 3$  and  $n_{\max} = 5$ .

Consider a state-space system with dimensions  $n_x, n_u \in \mathbb{N}$  corresponding to the number of states and control inputs, respectively. The vehicle's initial state is denoted  $z^0 \in \mathbb{R}^{n_x}$ , with  $n \in \mathbb{N}$  acquired target states  $z^j \in \mathbb{R}^{n_x}$  for  $j \in J$ , where  $J \subset \mathbb{N}$  is a finite indexing set for identifying the  $n$  targets. Additionally, there exist  $n - 1$  branch points  $\{z_k\}_{k=1}^{n-1} \subset \mathbb{R}^{n_x}$  with associated time horizons  $\{\tau\}_{k=1}^{n-1} \subset \mathbb{R}$ , from which  $n - 1$  deferrable segments exist between each  $z_{k-1}$  and  $z_k$ , starting with  $z_0 = z^0$ , along with suboptimality tolerances  $\{\epsilon\}_{k=1}^n$  for each target. Each target also has a bounding circle associated with the safety map (see Section III) that is updated at each branch point with a radius  $R^j \in \mathbb{R}$ . These radii will be implicitly considered a (dynamically-updated) parameter of  $z^j$  to simplify notation going forward (along with the scoring parameters discussed in Section III-C).

DDTO also requires the specification of a target rejection order  $\{\lambda_k\}_{k=1}^{n-2} \subset \mathbb{N}$ , where identification tags from  $J$  are specified in the order of deferral. To facilitate this, a desirability score is assigned to each target and represented by the evaluation function  $g_{\text{des}} : \mathbb{R}^{n_x} \rightarrow [0, \infty)$ . A proposed heuristic for this evaluation function is a weighted sum  $g_{\text{des}}(z^j) = w_{\text{des}}^\top c_{\text{des}}$ , where  $c_{\text{des}} \in \mathbb{R}^5$  contains the scores related to point cloud density, proximity to vehicle, site clustering, network prediction uncertainty, and landing zone size, as described in Section III-C, and  $w_{\text{des}} \in \mathbb{R}^5$  contains corresponding tunable weights for each element of  $c_{\text{des}}$ .

The function  $g_{\text{rem}} : \mathbb{R}^{n_x} \rightarrow \{0, 1\}$  determines if a target should be removed or not, and returns **true** if the target bounding radii exceeds the minimum safe bounding radius, denoted  $g_{\text{rem}}(z^j) = R^j < R_{\text{min}}$ .

The switching decision at a branch point is represented by a function  $g_{\text{switch}} : \mathbb{R}^{n_x} \rightarrow \{0, 1\}$  and returns **true** if the desirability score associated with the target in consideration for deferral exceeds all other targets' scores, such that  $g_{\text{switch}}(z^j) = g_{\text{des}}(z^j) > \max(\{g_{\text{des}}(z^i)\}_{i \in \{J/j\}})$ .

Under certain conditions, it may be desirable to *lock* the guidance to the best available target  $z^{j_{\text{lock}}}$ , disabling ADAPTIVE-DDTO re-computations and applying control towards this target at the current state  $z_{\text{lock}}$  (updated implicitly) with current time  $\tau_{\text{lock}}$  until the terminal condition  $z^{j_{\text{lock}}}$  is reached. A reasonable locking condition based on a minimum ‘‘cutoff’’ altitude  $h_{\text{cutoff}}$  could be used, such that  $g_{\text{lock}} : \mathbb{R} \rightarrow \{0, 1\}$  is defined with  $g_{\text{lock}}(h) = h < h_{\text{cutoff}}$ . To index into the altitude of  $z_k$ , the vector  $e_h = [0, 0, 1, 0, 0, 0]^\top$  is used.

There are also several other useful functions and relations that will be considered for the ADAPTIVE-DDTO formulation:

- $h_{\text{acquire}} : \mathbb{N} \rightarrow \mathbb{R}^{n_x}$  calls from HALSS to acquire  $n_{\text{max}} - n$  additional targets, where  $n \geq n_{\text{min}}$  is the current number of targets.
- $h_{\text{update}} : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$  calls from HALSS to update target  $z^j$  and associated properties.
- $\text{bisect} : \text{PROBLEM} \rightarrow \mathbb{N}$  performs a minimum-feasibility bisection search across the time-of-flight  $t_f \in \mathbb{R}$  for a fixed time-of-flight problem (such as LCvXPROBLEM) to return the optimal discretization node count  $N^*$  for a given discretization time step  $\Delta t \in \mathbb{R}$  (such that  $t_f^* = N^* \Delta t$ ).
- $\text{sort\_perm}_n : \mathbb{R} \rightarrow \mathbb{N}$  sorts the input collection of cardinality  $n \in \mathbb{N}$  in monotonically increasing order and returns the permutation indices that indicate the sorting order.
- $\text{comm\_control} : (\mathbb{R}^m \times \mathbb{N} \times \mathbb{R} \times \mathbb{R}) \rightarrow \emptyset$  extracts and commands the control sequence from DDTO-returned concatenated vector  $V \in \mathbb{R}^m$  (where  $m$  is the total number of control elements, see [11]) towards target  $j$  between time  $t_1$  and time  $t_2$ .

Algorithm 1 notably contains recursive calls to itself for *adaptive* re-computations when  $n < n_{\text{min}}$ .

## V. RESULTS

This section discusses implementation details and experimental results from an integrated simulation of a

---

### Algorithm 1: ADAPTIVE-DDTO

---

```

Input:  $z^0, J, \{z^j\}_{j \in J}$ 
/* Begin initialization */
1  $z_0 \leftarrow z^0$ 
2  $\{z^j\}_{j=1}^{n_{\text{max}}} \leftarrow h_{\text{acquire}}(n_{\text{max}} - |J|)$ 
3  $n \leftarrow n_{\text{max}}$ 
4  $J \leftarrow 1 : n$ 
5  $\{\lambda_k\}_{k=1}^{n-2} \leftarrow \text{sort\_perm}_n(\{g_{\text{des}}(z^j)\}_{j \in J})$ 
/* End initialization */
6 for  $j \in J$  do
7    $N^j \leftarrow \text{bisect}(\text{LCvXPROBLEM})$ 
8    $\{\tau_k\}_{k=1}^{n-1}, \{z_k\}_{k=1}^{n-1}, V \leftarrow$ 
     DDTO( $J, \{\lambda_k\}_{k=1}^{n-2}, z_0, \{z^j\}_{j \in J}, \{\epsilon^j\}_{j \in J}, \{N^j\}_{j \in J}$ )
9   for  $k = 1, \dots, n - 1$  do
10     $z^j \leftarrow h_{\text{update}}(z^j, k)$ 
11    if  $g_{\text{switch}}(z^j)$  then
12       $J \leftarrow \{J/k\}$ 
13    else
14       $J \leftarrow \{J/\lambda_k\}$ 
15    for  $j \in J$  do
16      if  $g_{\text{rem}}(z^j)$  then
17         $J \leftarrow \{J/j\}$ 
18     $n \leftarrow |J|$ 
19    if  $n < n_{\text{min}}$  then
20      ADAPTIVE-DDTO( $z_k, J, \{z^j\}_{j \in J}$ )
21    if  $g_{\text{lock}}(e_h^\top z_k)$  then
22       $j_{\text{lock}} \leftarrow \text{argmax}(\{g_{\text{des}}(z^j)\}_{j \in J})$ 
23       $\text{comm\_control}(V, j_{\text{lock}}, \tau_{k-1}, N^{j_{\text{lock}}} \Delta t)$ 
24    else
25       $\text{comm\_control}(V, \lambda_k, \tau_{k-1}, \tau_k)$ 

```

---

combined HALSS ADAPTIVE-DDTO implementation, as well as analytical results from a comparative study between ADAPTIVE-DDTO and DDTO.

#### A. Experimental Setup and Simulation Environment

An integrated implementation of HALSS and ADAPTIVE-DDTO for landing a quadrotor on a Martian surface (imaged by NASA HiRISE) was created using the Microsoft AirSim plugin to Unreal Engine 4, which provides a simulated multirotor air vehicle model. Notably, the simulation vehicle is exposed to Earth’s gravitational and atmospheric conditions for simplicity, however the methods proposed in this work can be extended without loss of generality to any environment and vehicle configuration (such as the Mars Helicopter [18]). The simulation algorithms are implemented in Python (HALSS) and Julia (ADAPTIVE-DDTO)<sup>2</sup>. Point clouds are generated using AirSim’s body-fixed LiDAR sensor model with a maximum range of 500 m and a circular scanning pattern with a 30° radial FOV.

<sup>2</sup>The code is available at <https://github.com/UW-ACL/HALO>

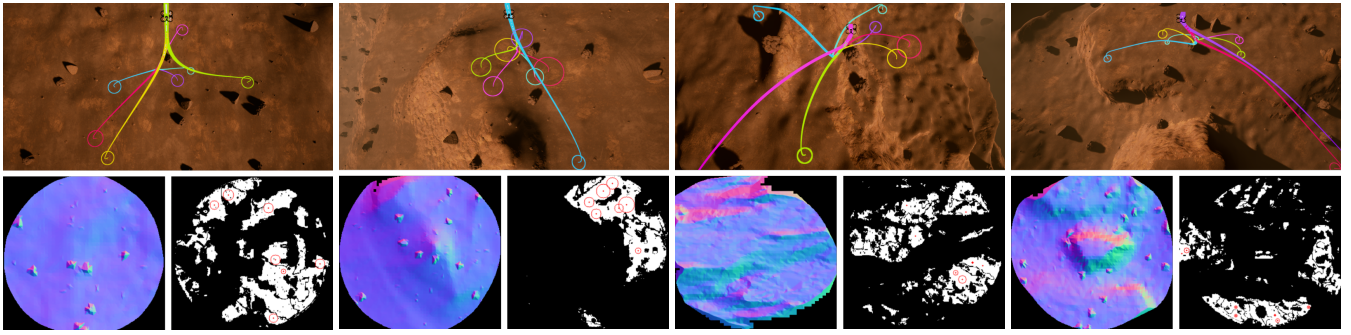


Fig. 7: Snapshots of four simulation instances on different terrain maps of increasing complexity from left to right. The top images shows our simulation environment during a run with the DDTO branches and landing sites, the bottom images show the global surface normal (left) and safety (right) maps with associated landing sites (red circles).

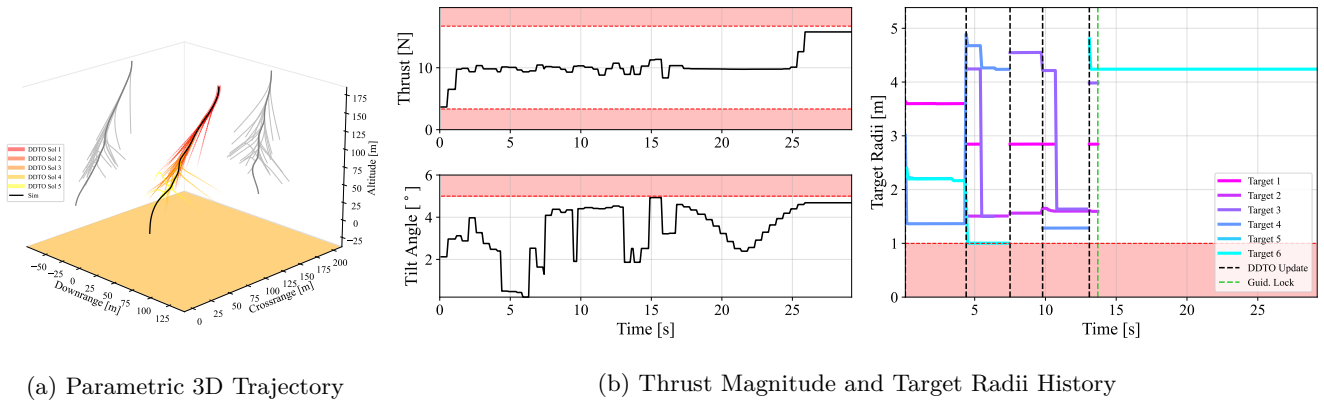


Fig. 8: Cumulative simulation results for the most complex (right-most) test case from Fig. 7.

### B. Integrated Simulation Results

Fig. 7 shows four different successful runs in four different environments of our coupled solution in the simulation environment<sup>3</sup>. Fig. 8a shows the full trajectory of the quadrotor as well as the branches of each DDTO solution in three dimensions. Notably, the initial state of the quadrotor for this simulation was above a large rock formation which caused the large divert. In Fig. 8b, to the left, we can see thrust and tilt angle control commands, which are shown to satisfy the optimization constraints (red regions). To the right, we can see the radius of each target as a function of time. Large vertical jumps in the radii correspond to previously unobserved rocks being identified as hazards within the radius of a given landing site. Smaller jumps can be attributed to hazards on the border of the circle becoming more defined and encroaching into the circle as the quadrotor descends. These results show that our coupled HALSS-ADAPTIVE-DDTO solution to the HALO problem is robust to larger divers as well as uncertainties in the landing sites.

### C. Comparative Study: DDTO vs. ADAPTIVE-DDTO

A comparative study was conducted between DDTO and ADAPTIVE-DDTO in a sandbox environment from a

<sup>3</sup>The following hyperparameters were used for all simulations:  $n_{\min} = 3$ ,  $n_{\max} = 7$ ,  $\epsilon^j = 0.1 \forall j \in J$ ,  $w_{\text{des}} = [0, 0, 1, 0, 0]^T$ ,  $h_{\text{cutoff}} = 65$  m,  $\alpha = 8^\circ$ ,  $\Delta t = 0.5$  s,  $T_{\min} = 3.34$  N,  $T_{\max} = 16.72$  N,  $\gamma_{\max} = 5^\circ$ ,  $v_{\max}^{\text{lat}} = v_{\max}^{\text{vert}} = 5$  m/s, and  $m = 1$  kg.

150 m initial altitude and targets generated in a uniform random distribution, where each target receives a 1% likelihood of failure every  $\Delta t$  seconds. Under 100 Monte Carlo trials, DDTO was observed to achieve successful landings 72.41% of the time with a mean cumulative control effort of 295.05 N-s, where as ADAPTIVE-DDTO achieved quantities of 85.56% and 293.72 N-s, respectively. This represents a 18.16% increase in terms of landing success-rate and 0.4% decrease in cumulative control effort from DDTO to ADAPTIVE-DDTO.

## VI. CONCLUSION

In this work, we addressed the Hazard-Aware Landing Optimization problem by introducing two novel frameworks, HALSS and ADAPTIVE-DDTO, which address the challenges of perception and planning, respectively. We demonstrate the effectiveness of our solutions in a realistic simulated environment and achieve near real-time performance for both HALSS and ADAPTIVE-DDTO. While this paper featured an integrated implementation, these algorithms can be used in stand-alone operations and are agnostic to the simulation environment. Moreover, our methods can apply to any application where an autonomous system has to navigate in an uncertain, dynamic, and potentially hazardous environment.

## REFERENCES

- [1] C. D. Epp, E. A. Robertson, and T. Brady, "Autonomous landing and hazard avoidance technology (alhat)," in *2008 IEEE Aerospace Conference*, 2008, pp. 1–7.
- [2] C. I. Restrepo, P.-T. Chen, R. R. Sostaric, and J. M. Carson, "Next-Generation NASA hazard detection system development," in *AIAA Scitech 2020 Forum*, ser. AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, Jan. 2020.
- [3] T. Ivanov, A. Huertas, and J. M. Carson, "Probabilistic hazard detection for autonomous safe landing," in *AIAA Guidance, Navigation, and Control (GNC) Conference*, ser. Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics, Aug. 2013.
- [4] R. Brockers, J. Delaune, P. Proença, P. Schoppmann, M. Dominik, G. Kubiak, and T. Tzanetos, "Autonomous safe landing site detection for a future mars science helicopter," in *2021 IEEE Aerospace Conference (50100)*, 2021, pp. 1–8.
- [5] P. F. Proença, J. Delaune, Rol, and Brockers, "Optimizing terrain mapping and landing site detection for autonomous uavs," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 9668–9674.
- [6] K. Konstantinidis and R. Förstner, "Guidance, navigation, and control system for landing near plume source on encedaladus," *J. Guid. Control Dyn.*, vol. 43, no. 12, pp. 2289–2308, Dec. 2020.
- [7] R. Moghe and R. Zanetti, "A deep learning approach to hazard detection for autonomous lunar landing," *J. Astronaut. Sci.*, vol. 67, no. 4, pp. 1811–1830, Dec. 2020.
- [8] K. Tomita, K. Skinner, K. Iiyama, B. Jagatia, T. Nakagawa, and K. Ho, "Hazard detection algorithm for planetary landing using semantic segmentation," in *ASCEND 2020*, ser. ASCEND. American Institute of Aeronautics and Astronautics, Nov. 2020.
- [9] K. Tomita, K. A. Skinner, and K. Ho, "Bayesian deep learning for segmentation for autonomous safe planetary landing," *arXiv:2102.10545*, Feb. 2021.
- [10] W. B. Greer and C. Sultan, "Shrinking horizon model predictive control method for helicopter–ship touchdown," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 5, pp. 884–900, 2020.
- [11] P. Elango, S. B. Sarsılmaz, and B. Açıkmeşe, "Deferring Decision in Multi-target Trajectory Optimization," in *AIAA SciTech 2022 Forum*, 2022, p. 1583.
- [12] V. B. Alex Kendall and R. Cipolla, "Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding," in *Proceedings of the British Machine Vision Conference (BMVC)*, G. B. Tae-Kyun Kim, Stefanos Zafeiriou and K. Mikolajczyk, Eds. BMVA Press, September 2017, pp. 57.1–57.12. [Online]. Available: <https://dx.doi.org/10.5244/C.31.57>
- [13] H. I. Choi, S. W. Choi, and H. P. Moon, "Mathematical theory of medial axis transform," *Pacific J. Math.*, vol. 181, no. 1, pp. 57–88, Nov. 1997.
- [14] B. Açıkmeşe and L. Blackmore, "Lossless convexification of a class of optimal control problems with non-convex control constraints," *Automatica*, vol. 47, no. 2, pp. 341–347, 2011.
- [15] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Acikmese, "Convex optimization for trajectory generation," *arXiv preprint arXiv:2106.09125*, 2021.
- [16] M. Fliess, J. Lévine, P. Martin, and P. Rouchon, "Flatness and defect of non-linear systems: introductory theory and examples," *International journal of control*, vol. 61, no. 6, pp. 1327–1361, 1995.
- [17] N. T. Nguyen, I. Prodan, F. Stoican, and L. Lefevre, "Reliable nonlinear control for quadcopter trajectory tracking through differential flatness," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6971–6976, 2017.
- [18] B. Balaram, T. Canham, C. Duncan, H. F. Grip, W. Johnson, J. Maki, A. Quon, R. Stern, and D. Zhu, "Mars helicopter technology demonstrator," in *2018 AIAA Atmospheric Flight Mechanics Conference*, 2018, p. 0023.