

# DEEP REINFORCEMENT LEARNING FOR SAFE LANDING SITE SELECTION WITH CONCURRENT CONSIDERATION OF DIVERT MANEUVERS \*

Keidai Iiyama<sup>†</sup>, Kento Tomita<sup>‡</sup>, Bhavi A. Jagatia<sup>§</sup>, Tatsuaki Nakagawa<sup>¶</sup>, and Koki Ho<sup>||</sup>

## ABSTRACT

This research proposes a new integrated framework for identifying safe landing locations and planning in-flight divert maneuvers. The state-of-the-art algorithms for landing zone selection utilize local terrain features such as slopes and roughness to judge the safety and priority of the landing point. However, when there are additional chances of observation and diverting in the future, these algorithms are not able to evaluate the safety of the decision itself to target the selected landing point considering the overall descent trajectory. In response to this challenge, we propose a reinforcement learning framework that optimizes a landing site selection strategy concurrently with a guidance and control strategy to the target landing site. The trained agent could evaluate and select landing sites with explicit consideration of the terrain features, quality of future observations, and control to achieve a safe and efficient landing trajectory at a system-level. The proposed framework was able to achieve 94.8 % of successful landing in highly challenging landing sites where over 80% of the area around the initial target landing point is hazardous, by effectively updating the target landing site and feedback control gain during descent.

## INTRODUCTION

On-board hazard detection and avoidance (HDA) capabilities are essential to enable new mission concepts that involve planetary surface operations. With a quick assessment of the perceived terrain data (e.g., DEM, visible spectrum map, or a combination thereof) from optical and/or LIDAR sensors, the HDA technology creates a map of probability of safety for prioritizing candidate landing zones. NASA has been actively developing the technology for Precise landing and Hazard Avoidance (PL& HA),<sup>1-6</sup> and Safe and Precise Landing Integrated Capabilities Evolution (SPLICE) project is currently underway to develop next generation PL&HA technologies.<sup>7-9</sup> The HDA algorithm being devel-

\*This paper is an updated version of Paper AAS 20-583 presented at the AAS/AIAA Astrodynamics Specialist Conference, Online, in 2020.

<sup>†</sup>Master Student, Department of Aeronautics and Astronautics, The University of Tokyo, 7-3-1, Hongo, Bunkyo-ku, Tokyo, Japan.

<sup>‡</sup>Ph.D. Student, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 30332.

<sup>§</sup>Master Student, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 30332.

<sup>¶</sup>Bachelor Student, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 30332.

<sup>||</sup>Assistant Professor, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 30332.




Figure 1: Landing site selection and control

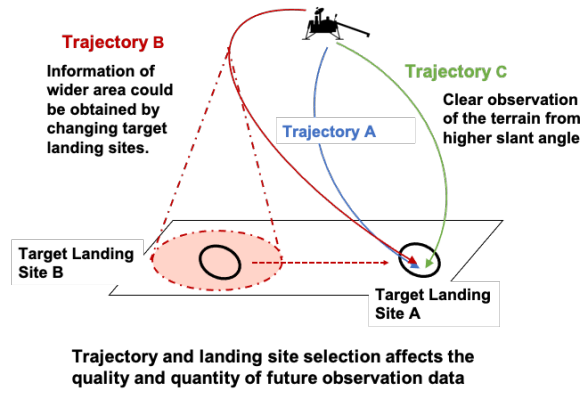


Figure 2: Observability and trajectory

oped for SPLICE directly leverages the HDA algorithms from the previous Autonomous Landing Hazard Avoidance Technology (ALHAT) project. ALHAT is capable of quickly assessing the DEM on-board and in real time during the descent, and assigns a probability of a safe landing to each pixel on the map.<sup>10-20</sup> ALHAT's HDA capability along with the terrain relative navigation and hazard detection functions was successfully demonstrated during the hardware-in-the-loop testing on the Morpheus Vertical Testbed.<sup>16,21</sup> However, a critical caveat of ALHAT is that it only selects a list of discrete landing points based on local static information (e.g., slopes, roughness, etc.). To consider the feasibility of the divert maneuver, a landing selection algorithm that calculates approximate landing footprint has been presented.<sup>22,23</sup> To reflect predetermined scientific values of each landing site in the landing site selection process, landing site selection methods that leverages Bayesian networks has been proposed.<sup>6</sup> Cui, et al. (2017) proposed a method that calculates synthetic landing area assessment criteria based on terrain safety, fuel consumption, and touchdown performance during descent.<sup>24</sup>

In these previously proposed landing site assessment methods, future changes in the target landing site are not considered. However, since the field of view and the quality of the observation data changes depending on the spacecraft state, the best landing site changes each time new observation data is obtained. This effect could not be ignored, especially when observability is limited or safe landing sites are sparse. When multiple chances of observation and divert maneuver are considered, the following two aspects should be considered.

The first aspect is the coupling between the target landing site selection and control maneuver planning. Existing landing site selection algorithms assume that the lander is guided by a simple predetermined control law, while the control law is designed to guide the lander to the decided landing target. However, when additional divert maneuvers are available in the future, the landing site selection policy is dependent on the control policy, and vice-versa. Therefore, the landing site and control plan has to be decided simultaneously in order to achieve an overall optimal trajectory. Fig.1 shows an example of this coupling. Suppose landing site A and landing site B both have uncertainty in the safety level that could be judged from the observation. If there is little residual fuel, it could be safer to take the

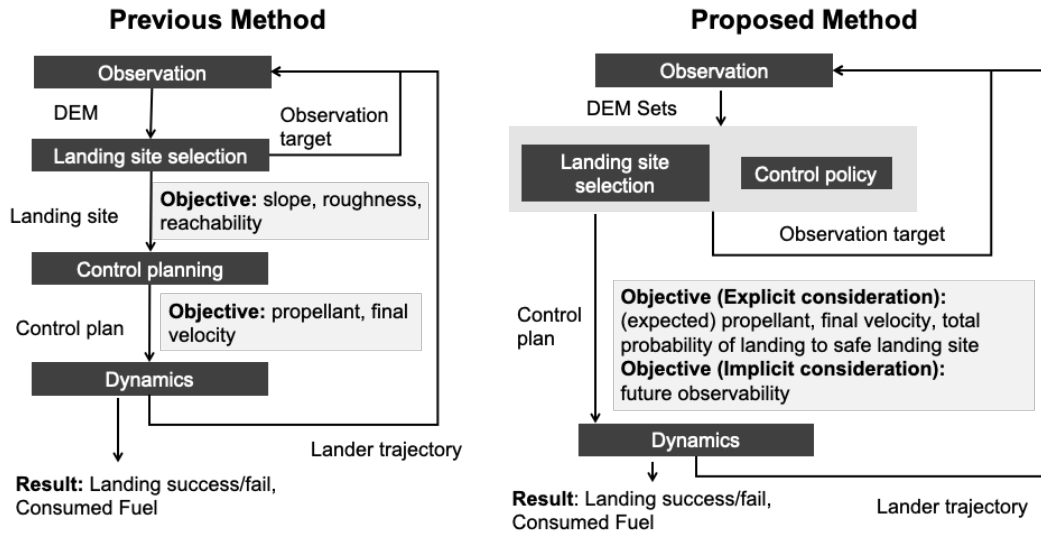


Figure 3: General Concept

fuel optimal path to the closer landing site B (Trajectory B). However, when there is sufficient residual fuel, it could be safer to take a longer path to landing site B (Trajectory A), if additional diverting to landing site A could happen in case landing site B turns out to be unsafe as the lander approaches the ground.

The second aspect is the observability of the terrain during descent. As illustrated in Figure 2, the change in the control policy and the target landing site affects the quality and number of the obtained observation data during descent. Therefore, the controller has to be aware that sufficient information to judge a good landing site could be obtained by following the planned trajectory. In general landing trajectory design, glide slope constraints are applied as a path constraint to ensure observations from higher elevations. However, the actual quality and quantity of observation information could not be explained only by slope angles. To tackle this problem, Crane (2013) developed an on-line information-seeking trajectory modification method that selects a trajectory that minimizes the weighted sum of the estimated entropy and field of view of the image obtained in the future, fuel consumption, and the coverage of the entire field.<sup>25</sup> However, the research focused on modifying the nominal trajectory for a fixed landing site, and updating the landing site successively during descent was not considered in the research. The impact of landing site selection on the observation data quality should also be considered, since changing the landing site affects the observation through changes in the trajectory and the observation target,

In order to tackle these problems, this paper proposes a learning-based method that selects landing site and design guidance trajectory successively during the HDA phase. The agent in the proposed method seeks to maximize the total probability of landing to a safe terrain landing site while minimizing total fuel consumption, landing error to the target, and final velocity. Model-free reinforcement learning techniques are leveraged to learn a policy that concurrently selects landing site and guidance strategy to the target, by interacting with the simulator environment and learning

how to maximize the total probability of successful landing with minimal fuel consumption. The general concept is shown in Fig.3.

The outline is as follows. First, we explain the assumptions used in modeling the HDA phase and show that the sequence of action choices can be formulated as a partially observable Markov decision process (POMDP). Next, the method to optimize landing site selection policies and the controller is introduced. Finally, the performance of the obtained controller is analyzed, and qualitative interception is given.

## PROBLEM FORMULATION

### Dynamics

A soft lunar landing scenario is assumed in this paper. In this paper, the 3 degrees of freedom (3-DOF) problem is considered. The equation of motion governing the dynamics of the problem are given as follows.

$$\begin{aligned} \mathbf{v} &= \dot{\mathbf{r}} \\ \mathbf{a} &= \frac{\mathbf{T}}{m} + \mathbf{g} \\ \dot{m} &= -\frac{\|\mathbf{T}\|}{g_{ref} I_{sp}} \end{aligned} \quad (1)$$

where  $\mathbf{r} = [r_x \ r_y \ r_z]^T$  is the position in the target centered orthonormal frame with the z-axis pointing upward, and  $\mathbf{g} = [0 \ 0 \ -1.62]^T$  is the gravity. In this paper, gravity is considered to be constant during the entire mission, and the effect of planetary rotation is ignored. In addition, limitation in the thrust range is applied as follows.

$$0 \leq \|\mathbf{T}\| \leq T_{max} \quad (2)$$

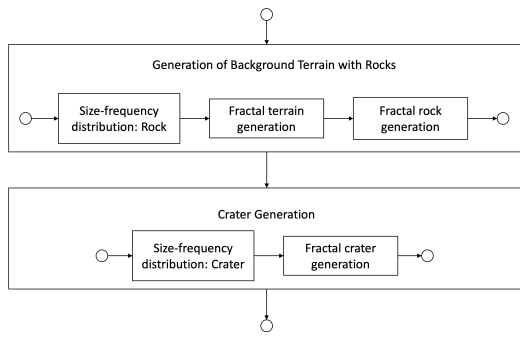
The following glide slope constraint are applied in most planetary landing problems

$$\theta_g = \arctan\left(\frac{\sqrt{r_x^2 + r_y^2}}{r_z}\right) < \theta_{gmax} \quad (3)$$


This constraint is applied in order to avoid the lander from hitting the terrain at low altitudes, and to ensure that the inclination to the target is kept over a certain amount for navigation and hazard detection purposes.

### Terrain Generation

Observation data consists of DEM generated by HD LIDAR, and spacecraft state information (position, velocity, mass) with noise. In order to generate LIDAR DEMs, we need a true DEM as a reference. Lunar Reconnaissance Orbiter (LRO) database provides Digital Terrain Model (DTM) of the lunar surface.<sup>26</sup> However, the resolution of the DTM ( $\approx 50\text{m/pixel}$ ) in the database is not sufficient for the LIDAR DEM model generation. Therefore, we made



**Figure 4:** Terrain generation process



**Figure 5:** Example of the generated terrain

an original terrain generator leveraging crater distribution models proposed by Pike, and boulder distribution models proposed by Bernard<sup>27,28</sup>. While the model was originally developed for Mars, studies of rock density on the moon show that the model can be substituted for the lunar topography by adjusting the parameters. The generated terrain has a size of 1000m x 1000m and a resolution of 1m. The generation process of the terrain is described in Fig.4. An example of the generated terrain is shown in Fig.5.

### Safety Map Generation

In order to assess if the landing has succeeded in the simulator, the safety of each landing point in the generated terrain map has to be assessed. We evaluated the safety by applying the safety assessment algorithm developed in the ALHAT project to the entire generated terrain. The algorithm for assessing the deterministic safety value  $V_D \in \{0, 1\}$  is shown in Algorithm 1.

### Observation Data Generation

The control agent utilizes two observation data for control: estimated lander state (position, velocity, and mass) and 2d map of the terrain. For the lander state, we assumed that we have perfect information without errors. Incorporating navigation errors and hazard relative navigation algorithms in the framework is important future work. For the 2d map of the terrain, we considered two options. The first option is to directly pass the LIDAR DEM, while the second option is to pass the stochastic safety map obtained by applying Algorithm1 on the obtained LIDAR DEM. The first option requires the control agent to assess the safety of each landing site from scratch, while in the second option, the control agent has access to more direct information about the safety of each landing point based on roughness and slope values. In this paper, we chose the second approach as we failed to design a controller using the first approach.

The simulation of the LIDAR DEM during descent requires complex calculations. To simulate a LIDAR DEM, detector pattern formulation, ray interception with the terrain, addition of the range bias, transformation to the point clouds, and transformation to a digital elevation map has to be conducted even if we assume that the LIDAR position is fixed to a known position. In lunar descent cases, errors from vehicle state knowledge, LIDAR misalignment, map

---

**Algorithm 1** Deterministic and Stochastic safety map generation

---

**Input:** DEM  $D(m \times n)$ **Output:** deterministic safety value map  $V_D(m \times n) \in \{0, 1\}$ , stochastic safety value map  $V_P(m \times n) \in [0, 1]$ **Initialize:** $S \leftarrow [0, \dots, 0]$  $\triangleright$  list of slope for each orientation $R \leftarrow [0, \dots, 0]$  $\triangleright$  list of roughness for each orientation $P \leftarrow [0, \dots, 0]$  $\triangleright$  list of safety probability for each orientation $O \leftarrow [0, \frac{1}{n_o}\pi, \dots, \frac{n_o-1}{n_o}\pi]$  $\triangleright$  possible lander footpad orientation**for**  $row = 1, \dots, m$  **do** $\triangleright$  for each pixel in the DEM**for**  $col = 1, \dots, n$  **do**

Calculate pad placement and contact

**for**  $o_i = 1, \dots, n_o$  **do** $\triangleright$  calculate worst case slope for all orientations $S[o_i] = \text{GET\_SLOPE}(row, col, D)$ **end for** $s_{max} \leftarrow \text{MAX}(S)$ **if**  $s_{max} > \text{safe slope threshold}$  **then** $\triangleright$  slope not safe $V[row][col] \leftarrow 0$ **else****for**  $o_i = 1, \dots, n_o$  **do** $R[o_i] = \text{GET\_ROUGHNESS}(row, col, D)$  $P[o_i] = \text{ROUGHNESS\_TO\_SAFETY\_PROBABILITY}(row, col, D, R)$ **end for** $V_P[row][col] \leftarrow \text{MIN}(P)$  $r_{max} \leftarrow \text{MAX}(R)$ **if**  $r_{max} > \text{safe roughness threshold}$  **then** $\triangleright$  roughness not safe $V_D[row][col] \leftarrow 0$ **else** $\triangleright$  slope and roughness safe $V_D[row][col] \leftarrow 1$ **end if****end if****end for****end for****return**  $V_D.V_P$ 

---

assembly errors when considering vehicle motions, and system latency should also have to be considered. Since the accurate modeling of the entire procedure is extremely complicated, as an initial concept study, we chose to generate pseudo observation safety map at each observation timing by cutting out the portion of the stochastic safety map of the entire field, and adding noise to it. The stochastic safety map could be obtained by directly applying Algorithm1 to the entire terrain DEM. By this way, we are able to greatly save the computation time required for simulation since we can avoid generating DEMs and running Algorithm1 at each observation timing.

When generating the "observation safety map", three main relationship between the lander's state (or trajectory) and the LIDAR DEM (and the generated safety map) were modeled. The first effect is the field of view (FOV) of the LIDAR DEM. For simplicity, we assumed the obtained DEM is square, and its two axes is always aligned with the




Figure 6: Illustration of the field of view of the LIDAR DEM

$x, y$  coordinate of the map. The length of one side of the square  $w_x = w_y$  is calculated using the following equation.

$$w_x = r_s \tan \phi_l \quad (4)$$

where  $r_s$  is the slant range (distance between the Lander and the target position), and  $\phi_l$  is the field of view of the LIDAR which was set to 11.4 [deg].(Fig. 6) In reality, the DEM field of view is not rectangle, and the field of view differs between the horizontal and vertical direction of the lander, but we believe this assumption is fair enough for an initial concept study.

The second effect is the effect of the slant range. As the slant range increases, errors in the range measurement in the LIDAR DEM increases. This effect was simulated by adding each pixel a Gaussian noise with an error proportional to the distance to each pixel. In addition, as the slant range gets longer, the sampling distance increases, small boulders in the terrain will be overlooked. This effect was modeled by fixing the LIDAR DEM size to 64x64. The third effect is the effect of slant angle. Slant angle is the angle between the horizontal plane and beam direction. As the slant angle decreases, holes in the DEM appear behind surface hazards. Therefore, regions where safety value could not be calculated appear in the safety map which leads to fewer safe regions in the map.<sup>29</sup> In this paper, this effect was simply modeled by assigning unsafe labels to pixels with slant angles between the lander smaller than 70 degrees.

In our settings, it is assumed that LIDAR DEMs could be obtained every 5 seconds, and the DEM size was fixed to 64x64. The interval time is based on the maximum processing time of the ALHAT algorithm to process DEM and generate a safety map. The entire observation data generation process is described in Algorithm 2.

### Modeling the Problem as POMDP

The HDA sequence is modeled as a POMDP  $P = \langle \mathcal{S}, \mathcal{A}, T, R, \Omega, O \rangle$  as follows.

- State space  $s \in \mathcal{S}$ : Spacecraft state (position, velocity, mass), true DEM of the entire field

---

**Algorithm 2** Observation Data Generation at Each Timestep

---

**Input:** true probabilistic safety value map of the entire terrain  $V_p(1000 \times 1000)$

true position of the lander in a surface fixed frame  $X_L = [x_l, y_l, z_l]$

**Output:** noisy probabilistic safety value map  $O_P(64 \times 64) \in [0, 1]$


```
// calculate the FOV of the DEM and return 2D array of center position of each pixel in the
// observed DEM
 $X_D, Y_D, w_x, w_y = \text{CALC\_FOV}(X_L, \phi_l)$  ▷  $\phi_l$ : LIDAR FOV, Eq.4
for  $row = 1, \dots, m$  do ▷ for each pixel in the DEM
  for  $col = 1, \dots, n$  do
     $x_d, y_d = X_d[row][col], Y_d[row][col]$ 
     $\theta = \text{SLANT\_ANGLE}(x_d, y_d, X_L)$ 
    if  $\theta > 20^\circ$  then ▷ slant angle to the pixel is over threshold
       $O_P[row][col] = 0$ 
    else
      // calculate the safety probability of the pixel by finding nearest neighbor pixel in the
      // whole  $V_P$  map
       $v_p = \text{NEAREST\_NEIGHBOR}(x_d, y_d, V_P)$ 
       $r = \text{SLANT\_RANGE}(x_d, y_d, X_L)$  ▷ [m]
       $O_P[row][col] = \text{CLIP}(v_p + \mathcal{N}(0, \frac{r}{500} * 0.05), 0, 1)$  ▷ Add noise to the safety value
    end if
  end for
end for
return  $O_P, w_x, w_y$ 
```

---

- Action space  $a \in \mathcal{A}$ : Lander thrust output, target landing position (determines next target position for LIDAR measurement)
- Reward space  $r \in \mathcal{R}$ : Consumed fuel, the safety of the final landing point, final velocity
- State transition function  $T : S \times A \rightarrow \Pi(S)$ : Stochastic dynamics of the spacecraft.  $T(s, a, s')$  is the probability of moving to state  $s' \in S$  when the agent at state  $s \in S$  takes action  $a \in \mathcal{A}$ .
- Observation space  $\Omega$ : LIDAR DEM, predicted spacecraft state
- Observation function  $O : S \times A \rightarrow \Pi(\Omega)$ : LIDAR DEM generation models, spacecraft state observation model

Obtaining an optimal policy (a policy that could maximize expected reward) in POMDP is difficult due to the partial observability of the problem. In general, an agent requires the entire history of the observation and action pairs  $h_t = \{a_0, r_0, o_1, \dots, a_{t-1}, r_{t-1}, o_t\}$ . POMDP could be converted into a belief Markov decision process (belief MDP) by introducing belief state  $b$ . Belief state is a conditional probability function for  $s \in S$  given the history  $h_t$ . When state transition function  $T$  and observation function  $O$  is known, an optimal policy of the POMDP could be obtained by approximating the belief MDP and applying a value iteration method.





**Figure 7:** Original observation data (Up) and reconstructed observation data(Down) by the Autoencoder

In this paper, we seek to obtain the optimal policy of the above POMDP without modeling the  $T$  and  $O$  function, but by learning from transition data collected by interacting with the simulator. There are two approaches in model-free approach for POMDPs. The first approach is the memory-less approach, which learns Markov policy by simply considering the most recent observation  $o$  as a state. In POMDP, the Bellman equation is not strictly satisfied, so deterministic policies leveraging only current state information is not guaranteed to be optimal. However, when the partial observability of the problem is weak, memory-less approach may be sufficient. The second approach is the memory-based approach, which learns history-dependent policy that uses the entire history data. This is usually achieved by using Recurrent Neural Networks (RNNs) that could store history information as a single state. In this paper, both memory-less and memory-based approaches are tested.

## **GUIDANCE AND CONTROL**

### **Observation Data Interpretation**

Since  $64 \times 64$  map data is used as part of the observation, policies that take in full image data taken as an input have too large dimension to optimize. Therefore, the whole history data is transferred into a low-dimension internal state value by leveraging an auto-encoder. The role of the auto-encoder is to extract an abstract, compressed representation of the LIDAR DEM data as a latent vector  $z$ , to avoid directly passing down large input data to the reinforcement learning agent. The auto-encoder was trained separately with the reinforcement learning agent, as proposed in the "World Models" paper by Ha (2018). The dataset for training was collected through 5000 random rollouts, and the auto-encoder was trained to minimize the difference between the observed safety map and the reconstructed safety map produced by the decoder. The following Fig.7 shows examples of the training data and the reconstructed image.

### **General Concept of Control**

For control, the most direct approach is to train the agent to output a three-dimensional thrust vector that guides the lander along a robust and fuel-efficient trajectory to the selected landing point. However, learning this feed-forward control policy from scratch requires high computational cost for learning. Several research have tackled this problem by shaping the reward in an effective way,<sup>30</sup> but we took a different approach in order to relate the control policy with target landing points. We adopted Zero-Effort-Miss/Zero Effort-Velocity (ZEM-ZEV) feedback

guidance algorithm as a baseline guidance law,<sup>31</sup> and designed the agent that controls the target landing point  $r_f$ , and adaptive hyper-parameters ( $K_R, K_V, t_{go}$ ) in the ZEM-ZEV guidance algorithm. The idea of learning adaptive hyper-parameters in ZEM-ZEV feedback guidance algorithm with reinforcement learning has already been proposed in previous research, to design an ZEM-ZEV feedback controller that avoids slant angle constraint violation during descent.<sup>32</sup> The difference in our paper is that the target position also changes during the descent.

## Controller

ZEM/ZEV feedback guidance algorithm calculates the optimal acceleration using the ZEM and ZEV, which represents the distance between the final target position and velocity and the projected final position and velocity when no additional control is added from current time  $t$ . The optimal acceleration is calculated as

$$\mathbf{a} = \frac{K_R}{t_{go}^2} \mathbf{ZEM} - \frac{K_V}{t_{go}} \mathbf{ZEV} \quad (5)$$

where  $K_R, K_V$  are control gains, and  $t_{go}$  is time-to-go. When there are no limitations in the thrust magnitude and no constraints in the trajectory, it is proved that energy-optimal trajectory could be obtained by setting the control gain as  $K_R = 6, K_V = 2$ . Energy optimal time-to-go  $t_{go}$  could be obtained by solving the following equation.

$$t_{go}^4 \mathbf{g}^T \mathbf{g} - 2t_{go}^2 (\mathbf{v}^T \mathbf{v} + \mathbf{v}_f^T \mathbf{v} + \mathbf{v}_f^T \mathbf{v}_f) + 12t_{go} (\mathbf{r}_f - \mathbf{r})^T (\mathbf{v} + \mathbf{v}_f) - 18(\mathbf{r}_f - \mathbf{r})^T (\mathbf{r}_f - \mathbf{r}) = 0 \quad (6)$$

When there is constraint  $|\mathbf{T}| = m|\mathbf{a}| \leq T_{max}$  in thrust magnitude, saturated acceleration is used for control as follows.

$$\mathbf{a} = \begin{cases} \bar{\mathbf{a}} & |\bar{\mathbf{a}}| \leq T_{max}/m \\ \bar{\mathbf{a}} T_{max}/m|\bar{\mathbf{a}}| & |\bar{\mathbf{a}}| > T_{max}/m \end{cases} \quad (7)$$

$$\bar{\mathbf{a}} = \frac{K_R}{t_{go}^2} \mathbf{ZEM} - \frac{K_V}{t_{go}} \mathbf{ZEV} \quad (8)$$

The goal of the research is to train a reinforcement learning agent that decides  $K_R, K_V, t_{go}$  and target landing position  $r_f$  at each time step.

## TRAINING

### Reinforcement Learning

In standard reinforcement learning problems, the agent interacts with the fully-observed environment  $E$ . At each time step  $t$ , the agent outputs an action  $a_t$ , and the environment returns the next observation  $o_{t+1} = s_{t+1}$  and a step reward  $r_{t+1} = g(s_t, a_t)$ . The agent decides its action based on its policy  $\pi(s_t)$ , which maps the state to a probability distribution over possible actions. The environment follows the transition dynamics  $p(s_{t+1}|s_t, a_t)$  and initial state

distribution  $p_{s_0}$ . The entire process could be modeled as a Markov decision process  $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, p_{s_0}, p, g\}$ . The goal of the agent is to learn a policy that maximizes the expectation of discounted future reward  $\mathbb{E}[R_t | \mathcal{M}, \pi]$  where  $R_t = \sum_{i=t}^T \gamma^T g(s_i, a_i)$ . Reinforcement learning algorithms utilize the following action-value function (or Q-function), which describes the expected return after taking action  $a_t$  in state  $s_t$ , then following policy  $\pi$ .

$$Q^\pi(s_t, a_t) = \mathbb{E}[R_t | s_t, a_t] \quad (9)$$

For the Q function, it is known that the following recursive equation called the Bellman equation holds.

$$Q^\pi(s_t, a_t) = \mathbb{E}_{r_t, s_{t+1} \sim E}[g(s_t, a_t) + \gamma \mathbb{E}_{a_{t+1} \sim \pi}[Q^\pi(s_{t+1}, a_{t+1})]] \quad (10)$$

### Deep Deterministic Policy Gradient (DDPG)

The Deep Deterministic Policy Gradient (DDPG) is an actor-critic, off-policy, model-free reinforcement learning algorithm for continuous action spaces.<sup>33</sup> In actor-critic methods, the critic estimates the action-value function  $Q(s, a)$ , while the actor produces the action given the current state based on its policy  $\pi(a|s)$ . As for the actor, we consider a parameterized deterministic policy  $a = \pi_\phi(s)$  with parameter  $\phi$ .  $\phi$  could be trained using the following deterministic policy gradient theorem and the estimated Q-values by the critic.<sup>34</sup>

$$\nabla_\phi J(\phi) = \mathbb{E}_{s \sim \rho_\pi}[\nabla_a Q(s, a|\theta)|_{s=s_t, a=\mu(s_t)} \nabla_\phi \mu(s|\phi)|_{s=s_t}] \quad (11)$$

For the critic, in order to handle continuous state and action spaces,  $Q(s, a)$  is also approximated using a function approximator parameterized by  $\theta$ . The critic uses the collected data to learn parameters that better approximate  $Q(s, a)$ . This could be achieved by minimizing the following loss  $L(\theta)$

$$L(\theta) = \mathbb{E}_{s_t \sim \rho^\beta, a_t \sim \beta, r_t \sim E}[Q(s_t, a_t|\theta) - y_t]^2 \quad (12)$$

where  $\beta$  is an arbitrary stochastic behavior policy, and  $\rho^\beta$  is a state visitation distribution using policy  $\beta$ .  $y_t$  is the target value

$$y_t = g(s_t, a_t) + \gamma Q(s_{t+1}, a_{t+1}|\theta) \quad (13)$$

In DDPG, deterministic policy is considered. For deterministic policy  $a = \mu(s)$ , the inner expectation in Eq.10 could be eliminated as follows.

$$Q^\mu(s_t, a_t) = \mathbb{E}_{r_t, s_{t+1} \sim E}[g(s_t, a_t) + \gamma Q^\mu(s_{t+1}, \mu(s_{t+1}))] \quad (14)$$

Therefore, when calculating the loss in Eq.12, transitions obtained from different stochastic policy could be used. The collected transition data  $(s_t, a_t, r_t, s_{t+1})$  is stored in a replay buffer, and then minibatch is sampled from the replay-buffer for loss calculation in Eq.12. This enables the samples for training to be distributed identically and independently which is the basic assumption for neural network training. When exploring the environment for collecting transition data, noise is added to the actor policy.

$$\bar{\mu}(s_t) = \mu(s_t, \theta_t) + \mathcal{N} \quad (15)$$

In addition, target networks are introduced in order to stabilize the training process and achieve greater convergence. When calculating Eq.13 with collected transition data  $(s_i, a_i, r_i, s_{i+1})$ , target critic with parameter  $\theta'$ , and target policy with parameter  $\phi'$  are used.

$$y_i = r_i + \gamma Q'(s_{i+1}, \mu'(s_{i+1}|\phi')|\theta') \quad (16)$$

Target networks are updated in order to slowly track the learned networks, as follows

$$\theta' \leftarrow \tau\theta + (1 - \tau)\theta' \quad (17)$$

$$\phi' \leftarrow \tau\phi + (1 - \tau)\phi' \quad (18)$$

where  $\tau \ll 1$ .

The key advantage of the DDPG method is that it enables to train deep reinforcement learning for continuous action space off-policy by assuming a deterministic policy. This greatly improves the sample efficiency because transition data obtained by different policies in the past could be re-used for training the agent.

### **Twin Delayed Deep Deterministic Policy Gradient (TD3)**

While DDPG has achieved significant performance in continuous control problems, the critical drawback is that it was sensitive to hyper parameter settings. This is because the overestimation of the Q-values makes the policy fall into local optima. The Twin Delayed Deep Deterministic Policy Gradient (TD3) algorithm made three major changes in the DDPG algorithm to overcome this drawback.<sup>35</sup> The first change (clipped double Q-learning) is to use two separate Q value approximators (critics), and use the smaller Q value of the two networks to calculate the target value  $y_t$ . The second change is target policy smoothing. For action in the target Q value, clipped noise is added to the actual action. This prevents the Q function from having sharp peaks by returning similar Q values for similar actions. With the clipped double Q-learning and target policy smoothing, the target value from transition data  $(s_i, a_i, r_i, s_{i+1})$  could be

calculated as follows.

$$y_i = r_i + \gamma \min_{i=1,2} Q'(s_{i+1}, \mu'(s_{i+1}|\phi') + \epsilon|\theta'_i), \quad \epsilon \sim \text{clip}(\mathcal{N}(0, \sigma), -c, c) \quad (19)$$

The last change is the delayed updates of the policy and the target networks. In order to stabilize the learning process, updates of the target networks and policy networks are carried out less frequently (example: once per every  $N$  steps) than the critic and networks.

### TD3 with memory

When the system is partially observed, optimal policy and action-value function both become function of the entire observation-action history  $h_t$ . Recurrent Deterministic Policy Gradient (RDPG) algorithm uses recurrent neural networks (RNN) in the policy and the action-value networks to preserve (limited) information of the history as part of the state.<sup>36</sup> In the RDPG setting,  $\mu(s)$  and  $Q(s, a)$  could be replaced with  $\mu(h)$  and  $Q(h, a)$ . Thus, the policy update could be obtained as follows.

$$\nabla_{\phi} J(\phi) = \mathbb{E}_{\tau \sim \nu_{\mu}} \left[ \sum_t \gamma^{t-1} \nabla_a Q(h, a|\theta)|_{h=h_t, a=\mu(h_t)} \nabla_{\phi} \mu(h|\phi)|_{h=h_t} \right] \quad (20)$$

The critic and actor networks are where  $\tau = (s_0, o_0, a_0, s_1, o_2, a_2, \dots)$  are drawn from the trajectory distribution  $\nu_{\mu}$  generated by the current policy  $\mu$ . In this paper, the RDPG algorithm was implemented by adding LSTM networks in the critic and actor networks. LSTM is a RNN architecture that is composed of a cell, an input gate, an output gate, and a forget gate. The critic and actor agent inputs the observation and action history  $h_t$  to their LSTM, and uses the output of the output gate  $\tilde{h}_t$  to calculate their outputs. The network structure for the critic and the actor is shown in Fig.8,9. This architecture is based on the works of Peng (2018) of robotic control. The upper network in the figure uses the current observation (and action for critic), while the bottom network utilizes past information stored in the LSTM. The network was trained with the TD3 algorithm, as shown in Algorithm 3. The difference from the original paper is that in our implementation, we created multiple replay buffers to store transition sequences of different episode length  $T$ . In our environment, the episode length varied from 4-10 time steps. We stored transition data with different length to separate replay-buffers so that we can create mini-batches with same episode length when training.

### Overall Framework and Training Process

The observation  $o_t$  which is used as an input of the agent is summarized in Table.1 and the output of the agent  $a_t$  is summarized in Table.2. For the training agent, we tested two architectures: a simple TD3 algorithm for memory-less approach, and a TD3 with LSTM approach for memory-based approach. For the architecture of TD3 agent without memory, we used the same architecture as the upper network of the policy and critic networks shown in Fig.8, 9. In




Figure 8: Architecture of the policy network




Figure 9: Architecture of the critic network

---

### Algorithm 3 TD3 with LSTM

---

Initialize Critic Network  $Q_{\theta_1}, Q_{\theta_2}$ , and actor network  $\mu_{\phi}$  with random parameters  $\theta'_1, \theta'_2, \mu'$

Initialize target networks  $\theta'_1 \leftarrow \theta_1, \theta'_2 \leftarrow \theta_2, \phi' \leftarrow \phi$

Initialize list of replay buffer list  $\mathcal{B}[]$

**for** episodes = 1, M **do**

    initialize empty history  $h_0$

$a_0 \leftarrow \text{RANDOM}(\mathcal{A}), t \leftarrow 0$

**while** not done **do**

$t \leftarrow t + 1$

        receive observation  $o_t$ , reward  $r_{t-1}$

$h_t \leftarrow h_{t-1}, a_{t-1}, o_t$

$\triangleright$  Append observation and previous action to history

        Select action  $a_t = \mu(h_t | \phi) + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma)$

        If done  $d_t \leftarrow 1$ , else  $d_t \leftarrow 0$

**end while**

    Store the sequence  $(a_0, o_1, r_0, d_0, \dots, a_{t-1}, o_t, r_t, d_t)$  in  $\mathcal{B}[t]$

$T \leftarrow \text{RANDOM}(t_{max})$

$\triangleright$  Select the sequence length to sample from

    Sample a mini-batch of  $N$  episodes with length  $T$ :

$(a_0^i, o_1^i, r_0^i, d_1^i, \dots, a_{T-1}^i, o_T^i, r_{T-1}^i, d_T^i)_{i=1, \dots, N}$  from  $\mathcal{B}[T]$

    Construct  $T \times N$  set of sequences  $h_t^i = (a_0^i, o_1^i, r_0^i, d_1^i, \dots, a_{t-1}^i, o_t^i, r_{t-1}^i, d_t^i)$

    Compute target values for each sample episode for  $t = 0, \dots, T - 1$

$$y_t^i = r_t^i + (1 - d_{t+1}^i) \gamma \min_{k=1,2} Q'(h_{t+1}^i, \mu'(h_{t+1}^i | \phi') + \epsilon | \theta'_k), \quad \epsilon \sim \text{clip}(\mathcal{N}(0, \sigma), -c, c)$$

    Update critic by minimizing the loss

$$\theta_k \leftarrow \text{argmin}_{\theta_k} \frac{1}{NT} \sum_i \sum_t (y_t^i - Q(h_t^i, a_t^i))^2$$

    Update actor parameter  $\phi$  by the deterministic policy gradient

$$\nabla_{\phi} J(\phi) = \frac{1}{NT} \sum_i \sum_t \nabla_a Q(h, a | \theta) |_{h=h_t^i, a=\mu(h_t^i)} \nabla_{\phi} \mu(h | \phi) |_{h=h_t^i}$$

**end for**

---

order to stabilize the training, all action parameters were scaled to  $[0,1]$  during training. In the actor, the target point at timestep  $t$  ( $= x_f^{(t)}, y_f^{(t)}$ ) is not outputted directly. Instead 2 variables  $\alpha_x, \alpha_y \in [-0.25, 0.25]$  is outputted from the

**Table 1:** Elements of the observed state  $o_t$ 

content	symbol	size
(true) lander position	$r_t$	3
(true) lander velocity	$v_t$	3
(true) lander mass	$m_t$	1
width of the observed DEM range in x,y direction	$w_x^{(t)}, w_y^{(t)}$	2
current target position	$x_f^{(t)}, y_f^{(t)}, z_f^{(t)}$	3
latent vector of the generated safety map encoded through autoencoder	$z_t$	32
total		44

**Table 2:** Elements of the action  $a_t$ 

content	symbol	size	range
gain of the ZEM-ZEV control	$K_R$	1	[5, 7]
gain of the ZEM-ZEV control	$K_V$	1	[1, 3]
degradation of time-to-go from previous timestep	$\delta t_{go}$	1	[4.25, 5.75]
target landing point position within the DEM FOV	$\alpha_x, \alpha_y$	2	[-0.5, 0.5]
total		5	

actor, and calculated the next target point using the following equation.

$$x_f^{(t)} = x_f^{(t-1)} + \alpha_x w_x, \quad y_f^{(t)} = y_f^{(t-1)} + \alpha_y w_y \quad (21)$$

In this way, the next target landing point is always selected to be within the FOV of the latest observed DEM. This prevents control failures caused by large deviations of the target between observations. The range of  $\alpha_x$  and  $\alpha_y$  was set to  $\alpha_x, \alpha_y \in [-0.25, 0.25]$ . The calculated new target point will be the center of the FOV of the DEM in the next step. Therefore,  $\alpha_x$  and  $\alpha_y$  are important output that has impacts on both guidance and observation. In addition, the actor does not directly output  $t_{go}$  of the next step, but instead output the degrade of the time-to-go until the next observation timing which will be conducted 5 seconds later ( $= \delta t_{go}$ ). The initial  $t_{go}$  is calculated using the initial lander position and target landing site and Eq.6.

Reward setting is the key in reinforcement learning. The reward at time step  $t$  was implemented as follows.

$$r_t = \alpha_m \frac{m_{t-1} - m_t}{m_0 - m_{dry}} + \alpha_f (U(z_{max} - z_t) + U(m_{dry} - m_t)) + \alpha_v d_t (|\mathbf{v}_t|) + \alpha_r d_t (|\mathbf{r}_t - \mathbf{r}_f^{(t-1)}|) + \alpha_s d_t V_D(r_t) \quad (22)$$

where

$$\begin{aligned}
 U(x) &= \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{otherwise} \end{cases} \\
 d_t &= \begin{cases} 1 & \text{if episode done} \\ 0 & \text{otherwise} \end{cases} \\
 V_D(r_t) &= \begin{cases} 1 & \text{if } r_t \text{ is a safe landing point} \\ -1 & \text{if } r_t \text{ is not a safe landing point} \end{cases} \\
 z_{max} &= 50 \quad (\text{maximum height of the terrain})
 \end{aligned} \tag{23}$$

The term with  $\alpha_m$  is a penalty for fuel consumption. The term with  $\alpha_f$  is a penalty term for breaking the two critical path constraints: hitting the ground and running out of fuel. The term with  $\alpha_v$  is a penalty for velocity norm at final episode to ensure soft landing, and the term with  $\alpha_r$  is a penalty for not landing to the final target landing point. These two terms are required because ZEM-ZEV control with adaptive gain and  $t_{go}$  parameters are not guaranteed to achieve a pinpoint soft landing to the target. Setting the weights  $\alpha_m, \alpha_v, \alpha_r, \alpha_s$  is a difficult problem which depends not only on the mission designer focus but the stability of the training process. It is natural to assume that  $\alpha_f, \alpha_v$  and  $\alpha_s$  should take relatively large values because hitting the ground at a high rate of speed and landing in a danger zone could both lead to the immediate loss of the lander. We used  $\alpha_m = 1, \alpha_f = -10, \alpha_v = 0.1, \alpha_r = 0.01, \alpha_s = 1$  in our research. Note that slope angle constraints are not considered as penalty in the reward settings. This is because slope angle constraints are not directly related to the primary objective of the planetary landing, but rather conservative constraints to assure preferable geometry for observation. By interacting with the simulator that simulates the observation process and the dynamics, we expect the agent to learn to take a trajectory with adequate slope angle in order to achieve safe landing. The overall framework is summarized in Fig 10.

## SIMULATION

### Initial Condition and Hyper-parameter Settings

The initial state and other specs of the lander are summarized in Table.3. Parameters that have range are randomly initialized within the range every time the episode starts. The true DTM is picked randomly from the list of DTMs generated before the training.

### Training Process

Fig.11 illustrates the transition of the training loss of the critic and actor network, reward, and the ratio of landing to safe landing site during training, for both the agent without LSTM and with LSTM. Values in the graph represent the



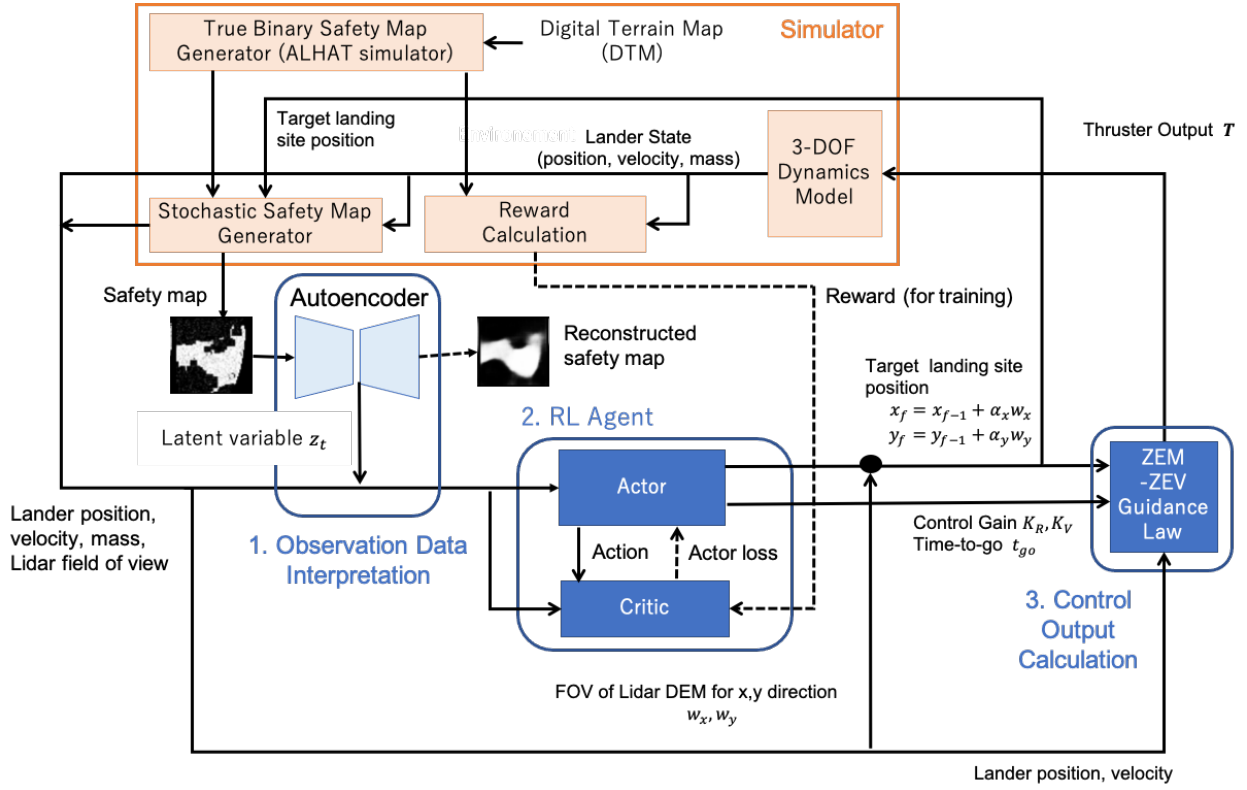



Figure 10: Overall framework

Table 3: Simulation condition

content	symbol	value
Initial altitude (downrange) [m]	$z_0$	[900, 1000]
Initial crossrange distance [m]	$\sqrt{x_0^2 + y_0^2}$	[200 250]
Initial downrange velocity [m/s]	$ v_z $	[20, 35]
Initial crossrange velocity [m/s]	$ v_z $	[5, 10]
Initial mass of the lander [kg]	$m_0$	1200
Altitude of the final target point [m]	$z_f$	50
Dry mass of the lander [kg]	$m_{dry}$	1150
Specific Impulse of the lander [s]	$I_{sp}$	325
Maximum thrust of the lander [N]	$T_{max}$	12000
Thrust magnitude error ratio	$ T - T_{plan} /T_{plan}$	0.05
Number of different DTMs used for training	$N_{D-training}$	70
Number of different DTMs used for variation	$N_{D-variation}$	18

moving average of 5000 episodes. When validating the performance during training, DTM datasets that are different from training datasets were used. It required around 250k iteration of training for both agents to reach around the maximum performance.

Contrary to our expectations, the agent without LSTM showed better performance compared with the agent with LSTM. The training of agent with LSTM was much more sensitive to its hyper-parameters, and we did not manage



**Figure 11:** Training log. It should be noted that the x-axis of the above two figures represent the number of training updates, while the x-axis of the bottom two figures represent the number of training episodes.

to achieve higher performance than TD3 without memory. Another reason why agent without LSTM could achieve high performance is likely because the partial observability of our problem is not strong, thanks to the access to the stochastic safety map with relatively low errors. Adding larger errors into the observation data such as navigation errors might require the use of memory to cope with strong partial observability.

### Comparison with other strategies

In order to assess the performance of the trained agent, the performance of the two agents were compared with two other strategies. The first strategy is the "Fixed Control Policy". This policy ignores the  $K_R, K_V, t_{go}$  output from the agent, and uses the agent only to select the target point. The gains are fixed to  $K_R = 6, K_V = 2$ , and  $t_{go}$  was calculated by Eq.6. This policy was introduced to see how the agent's adjustment of gain and flight time affected the trajectory, final velocity, and landing errors to the target. The second strategy is the "Single Divert Policy". This strategy makes no use of trained agents. This policy targets the initial target point until the altitude reaches 500m, and then selects the pixel with maximum safety value in the stochastic safety map obtained at 500m altitude as the final target point. The gains are fixed to  $K_R = 6, K_V = 2$ , and  $t_{go}$  was calculated by Eq.6. This policy was introduced to represent a simplified version of existing landing techniques.

The performance of the four policies was tested by landing simulation of 1000 episodes with random initial con-

**Table 4:** Comparison of performance between methods (1000 episodes. Mean value)

Criteria	Memory-Less Agent	Memory-Based	Fixed Control	Single Divert
Safe Landing Ratio [%]	94.8 %	88.3 %	93.2 %	89.7 %
Distance to final target [m]	2.492	4.547	3.786	0.005
Final Velocity [m/s]	0.637	1.520	1.189	0.557
Propellant Consumption [%]	69.88	63.70	67.84	68.89
Minimum Slant Angle [deg]	72.429	72.570	72.795	70.85
Maximum Thrust [N]	5818	6863	4553	4556


ditions. In order to assess the performance in cases that the agent needs to change its initial landing site, the initial target position was chosen so that at least 80% of the area within 100 meter radius is hazardous. Fig.12 shows the histogram of the reward, landing error to the final target point[m], final velocity magnitude[m/s], maximum thrust[N], and minimum slant angle [deg]. The mean value is summarized in Table.4.

The trained memory-less agent achieved maximum ratio (94.8%) of landing to safe landing site. The fixed control policy had a lower safe landing ratio to the agent policy, due to its limitation in changing the thrust magnitude flexibly when long divert maneuver is required. Fig.13 shows an example of such a situation. In this case, the initial target point (the center of the red square) is inside the hazardous region (black area of the figure), and the target landing point has to be shifted toward upside of the figure in order to achieve a safe landing. While the trained agent reduces the  $z$  axis velocity immediately as shown in Fig.13-(b) in order to maintain the altitude and wide field of view for searching safe landing sites, the fixed controller does not aggressively thrust the propellant as shown in Fig.13-(d). Therefore the field of view of the fixed gain controller is kept limited, and the fixed controller fails in finding a wide safe area to land.

The trained agent with memory had the lowest probability of landing to safe landing site, and the largest landing position error and velocity error among the four policies. On the contrary, the average fuel consumption was the lowest among the four policies. This implies that the agent was trapped in a local optima policy that increases reward by reducing fuel consumption rather than improving other metrics.

The single observation policy fails in selecting safe landing sites when safe landing sites could not be obtained in the single observation due to limitations in the field of view or slant angles. The rate of failure may be reduced by selecting optimal altitude for the divert decision or implementing better landing site selection algorithms from safety maps. However, it should be noted that in reality, the risk of single observation policy is likely to increase than our simulator due to additional error sources (navigation errors, attitude constraints, etc) which are not simulated in our research. These errors degrade the quality of the DEM (or safety map) and the accuracy of lander guidance to the target point.

While the trained agent showed high potential in adaptly changing its target position and control outputs to find




**Figure 12:** Evaluation of performance. 'Agent' is the policy of the trained memory-less agent, 'LSTM' is the policy of the trained memory-based agent, 'Fixed' is the policy with fixed ZEM-ZEV parameters, and 'Single' is the policy that conducts single divert maneuver and selects the landing site by directly using the stochastic safety map.

safe landing sites, they also have limitations in reducing velocities and landing error to the final target point decided by the agent. Since convergence to the target position and velocity is not guaranteed when control gain in ZEM-ZEV control is changed, this disadvantage is inevitable without further refinements. In addition, fuel consumption of the memory-less trained agent was larger compared with the two agents for comparison, due to the large thrust magnitude during descent and frequent changes of the target point. Addressing the trade-off between safety and fuel consumption from the point of the frequency of the target point changes is an interesting future work.

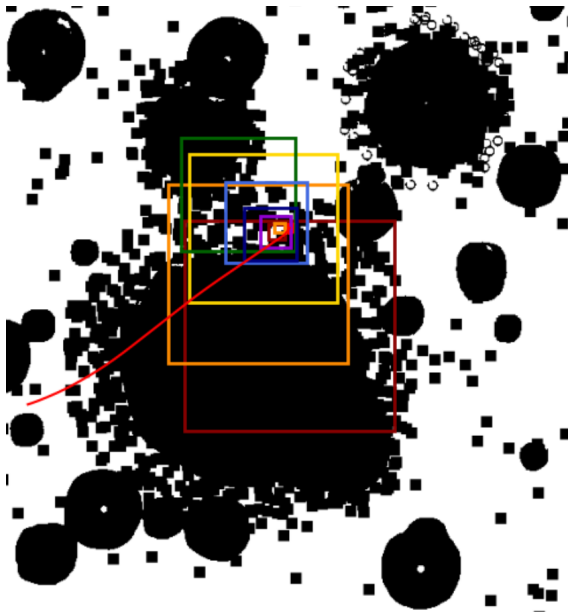
## CONCLUSION

In this paper, we proposed a new learning-based framework for Hazard Detection and Avoidance (HDA) phase which successively updates the target landing site and control parameters simultaneously after each observation, in order to cope with the coupling between observation, guidance, and control. We modeled the HDA sequence as a POMDP, and a reinforcement learning agent that interprets the obtained map using auto-encoder and outputs control parameters for ZEM-ZEV feedback control law was developed to find an optimal policy. The agent was trained by interacting with the simulator, and the trained agent was able to achieve over 90% probability of successful landing at difficult landing sites where over 80% of the terrain was hazardous around the initial target point, by gradually updating the target landing point towards safe regions.



In order to incorporate a more realistic and higher level of uncertainty to the environment, our future work will





(a) Trajectory by agent controller (White:Safe Black:Unsafe)



(c) Trajectory by fixed controller (White:Safe Black:Unsafe)




(b) Control history by the agent controller (Up:Velocity Down:Thrust)



(d) Control history by the fixed controller ((Up:Velocity Down:Thrust)

**Figure 13:** Comparison of the obtained agent and a ZEM-ZEV controller with fixed gain. Red line in figure(a),(c) shows the lander trajectory, while the square shows the field of view (FOV) of the LIDAR DEMs.

mainly focus on the refinement of the simulator model. This includes the incorporation of accurate LIDAR DEM generation models, expansion of dynamics to 6-DOF, and incorporation of lander state measurement errors. As the partial observability of the environment increases by incorporating various error sources, a memory-based agent might be required for sufficient performance. We are also planning to test other baseline control policies and observation



**Figure 14:** Changes in the control gain  $K_R$ ,  $K_V$  while control in Fig.13-(a),(b)

data interpretation architectures to improve the agent's performance and stability.

### ACKNOWLEDGEMENTS

This material is partially based upon work supported by the National Aeronautics and Space Administration under Grant No.80NSSC20K0064 through the NASA Early Career Faculty Program.

## REFERENCES

- [1] A. E. Johnson, A. R. Klumpp, J. B. Collier, and A. A. Wolf, "Lidar-based hazard avoidance for safe landing on Mars," *Journal of guidance, control, and dynamics*, Vol. 25, No. 6, 2002, pp. 1091–1099.
- [2] N. Serrano and H. Seraji, "Landing site selection using fuzzy rule-based reasoning," *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 4899–4904.
- [3] A. Huertas, Y. Cheng, and L. H. Matthies, "Real-time hazard detection for landers,," *NASA Science Technology Conference*, 2007.
- [4] L. Matthies, A. Huertas, Y. Cheng, and A. Johnson, "Stereo vision and shadow analysis for landing hazard detection," *2008 IEEE International Conference on Robotics and Automation*, 2008, pp. 2735–2742.
- [5] Y. Cheng, A. E. Johnson, L. H. Mattheis, and A. A. Wolf, "Passive imaging based hazard avoidance for spacecraft safe landing," *I-SIRAS 2001: 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2001, pp. 1–14.
- [6] N. Serrano, "A bayesian framework for landing site selection during autonomous spacecraft descent," *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 5112–5117.
- [7] C. I. Restrepo, R. Lovelace, R. R. Sostaric, J. M. Carson, and N. G. Spaceflight, "NASA SPLICE Project : Developing the Next Generation Hazard Detection System," *2nd RPI Space Imaging Workshop*, 2019, pp. 2–3.
- [8] A. D. Cianciolo, S. Striepe, J. Carson, R. Sostaric, D. Woffinden, C. Karlgaard, R. Lugo, R. Powell, and J. Tynis, "Defining navigation requirements for future precision lander missions," *AIAA Scitech 2019 Forum*, No. January, 2019, pp. 1–18, 10.2514/6.2019-0661.
- [9] J. M. Carson, M. M. Munk, R. R. Sostaric, J. N. Estes, F. Amzajerddian, J. Bryan Blair, D. K. Rutishauser, C. I. Restrepo, A. D. Cianciolo, G. T. Chen, and T. Tse, "The splice project: Continuing nasa development of gn&c technologies for safe and precise landing," *AIAA Scitech 2019 Forum*, No. January, 2019, pp. 1–9, 10.2514/6.2019-0660.
- [10] T. Brady, E. Robertson, S. P. C. App, and D. Zimpfer, "Hazard detection methods for lunar landing," *2009 IEE Aerospace Conference*, 2009, pp. 1–18.
- [11] C. D. App and T. B. Smith, "Autonomous precision landing and hazard detection and avoidance technology (ALHAT)," *2007 IEEE Aerospace Conference*, 2007, pp. 1–7.
- [12] T. Ivanov, A. Huertas, and J. M. Carson, "Probabilistic Hazard Detection for Autonomous Safe Landing," *AIAA Guidance, Navigation, and Control (GNC) Conference*, 2013, p. 5019.
- [13] T. Brady and J. Schwartz, "ALHAT system architecture and operational concept," *2007 IEEE Aerospace Conference*, 2007, pp. 1–13.
- [14] S. A. Striepe, C. D. Epp, and E. A. Robertson, "Autonomous precision landing and hazard avoidance technology (ALHAT) project status as of May 2010," *International Planetary Probe Workshop 2010 (IPPW-7)*, 2010.
- [15] D. Rutishauser, C. Epp, and E. Robertson, "Free-Flight Terrestrial Rocket Lander Demonstration for NASA's Autonomous Landing and Hazard Avoidance Technology (ALHAT) System," *AIAA SPACE 2012 Conference & Exposition*, 2012, p. 5239.
- [16] C. Epp, E. Robertson, and J. M. Carson, "Real-time hazard detection and avoidance demonstration for a planetary lander," *AIAA SPACE 2014 Conference and Exposition*, 2014, p. 4312.

- [17] A. E. Johnson, A. Huertas, R. A. Werner, and J. F. Montgomery, "Analysis of on-board hazard detection and avoidance for safe lunar landing," *2008 IEEE Aerospace Conference*, 2008, pp. 1–9.
- [18] A. Huertas, A. E. Johnson, R. A. Werner, and R. A. Maddock, "Performance evaluation of hazard detection and avoidance algorithms for safe Lunar landings," *2010 IEEE Aerospace Conference*, 2010, pp. 1–20.
- [19] B. E. Cohan and B. K. Collins, "Landing point designation algorithm for lunar landing," *Journal of Spacecraft and Rockets*, Vol. 46, No. 4, 2009, pp. 858–864.
- [20] S. Paschall and T. Brady, "Demonstration of a safe & precise planetary landing system on-board a terrestrial rocket," *2012 IEEE Aerospace Conference*, 2012, pp. 1–8.
- [21] N. Trawny, A. Huertas, M. E. Luna, C. Y. Villalpando, K. Martin, J. M. Carson, A. E. Johnson, C. Restrepo, and V. E. Roback, "Flight testing a Real-Time Hazard Detection System for Safe Lunar Landing on the Rocket-Powered Morpheus Vehicle," *AIAA Guidance, Navigation, and Control (GNC) Conference*, 2015.
- [22] N. Serrano, "A Bayesian framework for landing site selection during autonomous spacecraft descent," *IEEE International Conference on Intelligent Robots and Systems*, 2006, pp. 5112–5117, 10.1109/IROS.2006.282603.
- [23] S. R. Ploen, H. Seraji, and C. E. Kinney, "Determination of spacecraft landing footprint for safe planetary landing," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 45, No. 1, 2009, pp. 3–16, 10.1109/TAES.2009.4805259.
- [24] P. Cui, D. Ge, and A. Gao, "Optimal landing site selection based on safety index during planetary descent," *Acta Astronautica*, Vol. 132, No. July 2016, 2017, pp. 326–336, 10.1016/j.actaastro.2016.10.040.
- [25] E. S. Crane and S. M. Rock, "Guidance augmentation for reducing uncertainty in vision-based hazard mapping during lunar landing," *IEEE Aerospace Conference Proceedings*, 2013, 10.1109/AERO.2013.6496946.
- [26] M. Barker, E. Mazarico, G. Neumann, M. Zuber, J. Haruyama, and D. Smith, "A new lunar digital elevation model from the Lunar Orbiter Laser Altimeter and SELENE Terrain Camera," *Icarus*, Vol. 273, 2016, pp. 346 – 355, <https://doi.org/10.1016/j.icarus.2015.07.039>.
- [27] R. J. Pike, "Size-dependence in the shape of fresh impact craters on the moon.," *Impact and Explosion Cratering: Planetary and Terrestrial Implications* (D. J. Roddy, R. O. Pepin, and R. B. Merrill, eds.), Jan. 1977, pp. 489–509.
- [28] D. E. Bernard and M. P. Golombek, "Crater and rock hazard modeling for Mars landing," *AIAA Space 2001 Conference and Exposition*, Albuquerque, NM, 2001, 10.2514/6.2001-4697.
- [29] C. I. Restrepo, P.-T. Chen, R. R. Sostaric, and J. M. Carson, "Next-Generation NASA Hazard Detection System Development," 2020, pp. 1–10, 10.2514/6.2020-0368.
- [30] B. Gaudet, R. Linares, and R. Furfaro, "Deep Reinforcement Learning for Six Degree-of-Freedom Planetary Landing," *Advances in Space Research*, 01 2020, 10.1016/j.asr.2019.12.030.
- [31] B. Ebrahimi, M. Bahrami, and J. Roshanian, "Optimal sliding-mode guidance with terminal velocity constraint for fixed-interval propulsive maneuvers," *Acta Astronautica*, Vol. 62, No. 10, 2008, pp. 556 – 562, <https://doi.org/10.1016/j.actaastro.2008.02.002>.
- [32] R. Furfaro, A. Scorsoglio, R. Linares, and M. Massari, "Adaptive Generalized ZEM-ZEV Feedback Guidance for Planetary Landing via a Deep Reinforcement Learning Approach," 2020, 10.1016/j.actaastro.2020.02.051.



- [33] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *ICLR* (Y. Bengio and Y. LeCun, eds.), 2016.
- [34] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, “Deterministic Policy Gradient Algorithms,” *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*, ICML’14, JMLR.org, 2014, p. I–387–I–395.
- [35] S. Fujimoto, H. Hoof, and D. Meger, “Addressing Function Approximation Error in Actor-Critic Methods,” *International Conference on Machine Learning*, 2018, pp. 1582–1591.
- [36] N. M. O. Heess, J. J. Hunt, T. P. Lillicrap, and D. Silver, “Memory-based control with recurrent neural networks,” *ArXiv*, Vol. abs/1512.04455, 2015.