Bayesian Deep Learning for Segmentation for Autonomous Safe Planetary Landing *

Kento Tomita[†] Georgia Institute of Technology, Atlanta, GA, 30332

Katherine A. Skinner[‡] University of Michigan, Ann Arbor, MI, 48109

Koki Ho[§] Georgia Institute of Technology, Atlanta, GA, 30332

Hazard detection is critical for enabling autonomous landing on planetary surfaces. Current state-of-the-art methods leverage traditional computer vision approaches to automate the identification of safe terrain from input digital elevation models (DEMs). However, performance for these methods can degrade for input DEMs with increased sensor noise. In the last decade, deep learning techniques have been developed for various applications. Nevertheless, their applicability to safety-critical space missions has often been limited due to concerns regarding their outputs' reliability. In response to these limitations, this paper proposes an application of the Bayesian deep-learning segmentation method for hazard detection. The developed approach enables reliable, safe landing site detection by: (i) generating simultaneously a safety prediction map and its uncertainty map via Bayesian deep learning and semantic segmentation; and (ii) using the uncertainty map to filter out the uncertain pixels in the prediction map so that the safe site identification is performed only based on the certain pixels (i.e., pixels for which the model is certain about its safety prediction). Experiments are presented with simulated data based on a Mars HiRISE digital terrain model by varying uncertainty threshold and noise levels to demonstrate the performance of the proposed approach.

I. Introduction

A UTONOMOUS landing is identified as a medium-high priority for future NASA missions [1, 2]. Real-time hazard detection and avoidance (HDA) is a critical component for enabling autonomous landing. The current state-of-the-art method, developed for the Autonomous Landing Hazard Avoidance Technology (ALHAT) program, demonstrated potential for computer vision algorithms to automate online identification of safe or unsafe terrain from

^{*}This paper is a substantially revised version of Paper AAS 21-253 presented at the AAS/AIAA Space Flight Mechanics Meeting, Virtual, 2021. †Ph.D. Student, Daniel Guggenheim School of Aerospace Engineering, AIAA Student Member.

[‡]Assistant Professor, Department of Robotics, AIAA Member.

[§]Associate Professor, Daniel Guggenheim School of Aerospace Engineering, AIAA Senior Member.

input Digital Elevation Models (DEMs) [3–6], which can be generated from measurements from a flash Light Detection and Ranging (LiDAR) sensor during a mission. The Chang'e-3 lunar lander mission, the only space mission that has successfully demonstrated autonomous HDA, also explicitly evaluated hazards based on DEMs in the precise hazard avoidance phase [7]. However, DEM-based explicit evaluation of hazards is sensitive to severe sensor noise, navigation error, and missing data, which can lead to errors in safe landing site prediction.

Recent work has demonstrated the potential for deep neural networks (DNNs) to improve both accuracy and computational cost for onboard hazard detection [8, 9]. These prior works take input DEMs to output a binary map labeling each pixel as safe or unsafe for landing. However, one caveat in machine learning methods is their output's reliability. Most such methods rely on gathering a representative dataset for training, and may not generalize well to test data that is outside of the training distribution. In these cases, a predicted safety map will still be returned without any indication of how reliable that prediction is based on the input DEM.

There has been recent interest in estimating network uncertainty in deep learning predictions to provide further insight on the reliability of the network prediction [10–14]. However, these works have not been evaluated for autonomous landing on planetary surfaces, and there is a lack of literature on how to translate output uncertainty estimates into the safety map generation.

This work focuses on the development and validation of Bayesian deep learning methods for hazard detection for landing on a planetary surface. The specific contributions of this work are (i) to perform semantic segmentation and uncertainty prediction for identifying safe landing sites from input DEMs, (ii) to develop an algorithm that takes in the safety predictions and uncertainty maps to produce an uncertainty-aware (and thus reliable) identification of safe landing sites, and (iii) to perform experiments for demonstration of developed methods on DEMs by varying sensor noise and the uncertainty threshold.

II. Background

A. Autonomous Hazard Detection

The ALHAT program was launched by NASA in 2005 to develop advanced capabilities for spacecraft autonomous hazard detection and precise landing on lunar and planetary surfaces [15–17]. The work developed during the ALHAT program is considered state-of-the-art in autonomous hazard detection for planetary landing. In particular, Ivanov et al. [18] developed a probabilistic method for autonomous hazard detection that considers the vehicle size, the vehicle-surface configuration during landing, and navigation errors. The algorithm searches over the potential footpad positions and lander orientations during landing to determine the slope of the lander. The slopes are computed based on where the footpads fall on the terrain, as well as the roughness of hazards underneath the lander based on the final lander pose, i.e. position and orientation. If there are unsafe configurations for the lander at a given point on the DEM, this

point is labeled as unsafe for landing. This provides pixel-wise labeling of safe and unsafe landing points. The final output of the algorithm is a list of proposed safe landing sites, along with their coordinates and safety probability. This work was integrated on board the Morpheus vehicle and has been extensively tested in real flight scenarios. Flight tests demonstrated successful autonomous hazard detection and hazard relative navigation capabilities for the Morpheus vehicle over realistic lunar-like terrain [5, 6, 19]. With simplifying assumptions, the ALHAT algorithm developed by Ivanov et al. is capable of operating in real-time for efficient and effective hazard detection during landing. However, the performance of the algorithm can degrade with increased sensor noise, navigation error, and missing data. Building off of the success of the ALHAT project and Morpheus experiments, next-generation hazard detection and avoidance systems are currently under development through NASA's Safe & Precise Landing - Integrated Capabilities Evolution (SPLICE) program [20].

B. Machine Learning for Safe Landing on Planetary Surfaces

Recent work has focused on developing machine learning-based approaches to aid in guidance, navigation, and control for spacecraft landing on planetary surfaces [21–24]. For scenarios where high-resolution maps are available prior to landing, these methods can learn guidance and control policies to target the selected landing site. Our work focuses on the perceptual task of hazard detection to determine safe and unsafe landing sites when the DEM is computed onboard to provide additional information to the guidance system. This capability is critical for applications where high-resolution maps are not available prior to landing, and for detection of small hazards such as rocks.

There are several prior works that leverage machine learning for detecting and labeling geological features such as rocks and craters on the surface of the moon [25–30], Mars [31], or small bodies [32]. Our proposed work instead focuses specifically on detecting safe sites for landing without explicitly detecting hazardous terrain features such as craters or rocks. Recently, Moghe and Zanetti [9] developed a DNN for detecting safe and unsafe areas in a DEM for automated hazard detection. Their work develops a novel loss function specifically designed to decrease the false positive (safe) rate, since false positive cases indicate that an area is predicted safe when it is actually unsafe, which can lead to catastrophic situations for autonomous landing operations. Tomita et al. [8] also explored several DNN architectures for real-time hazard detection. In particular, their results demonstrated that a semantic segmentation DNN can outperform the current state-of-the-art ALHAT algorithm for the prediction of safe pixels when the input DEM has sensor noise. Our work builds on these prior works to include estimation of network uncertainty for the output binary safety map to provide a cue on the reliability of the network prediction. Our proposed approach leverages output uncertainty maps to enable uncertainty-aware safe site identification. Furthermore, we provide experiments to evaluate our approach with varying noise parameters for testing the network.

C. Uncertainty Estimation

There are two main types of uncertainty that can be modeled: aleatoric uncertainty and epistemic uncertainty [33]. *Aleatoric uncertainty* is uncertainty due to sensor noise or motion noise [33], whereas *Epistemic uncertainty*, which is also referred to as model uncertainty, includes uncertainty in the model structure and parameters [33]. Estimating uncertainty through deep learning frameworks has been investigated in the literature [34]. For learning-based frameworks, aleatoric uncertainty is the main source of uncertainty if the test distribution is well-aligned with the training distribution (which is a common assumption in classical machine learning problems). However, predictions with low aleatoric uncertainty may still be incorrect due to the epistemic uncertainty (e.g., uncertainty in the DNN parameters), especially if the test distribution contains samples that fall outside of the training distribution. Particularly, epistemic uncertainty is an important contributor to the overall uncertainty when training sets are small or limited in variability, or for safety-critical applications. Epistemic uncertainty can be estimated through approximate Bayesian inference techniques, such as ensembling [35] or Monte Carlo (MC) dropout [34]. For the application of autonomous landing on planetary surfaces, it is critical to account for both aleatoric and epistemic uncertainty to ensure that the safety map estimate considers uncertainty due to sensor noise and uncertainty due to encountering test cases that fall outside of the training distribution. Thus, our work will leverage prior work to estimate aleatoric and epistemic uncertainty, in order to enable uncertainty-aware safety map generation for a planetary lander.

III. Proposed Methodology

Figure 1 shows an overview of our developed pipeline for learning-based safety map generation. The pipeline takes a noisy DEM obtained by noisy LiDAR scans as input to the semantic segmentation network. We assume that the raw output from LiDAR is converted to the DEM with associated information of the lander's position and attitude. The base network is the Bayesian semantic segmentation network, Bayesian SegNet [34]. Bayesian SegNet leverages MC-dropout in which dropout is included during both training and testing. During inference, this allows for an output of both a pixel-wise safety prediction and an associated uncertainty that indicates how certain the network is in the prediction. The safety map and the associated uncertainty map are used to compute the uncertainty-aware safety map.

A. Bayesian Deep Learning for Semantic Segmentation

Figure 2 shows the network architecture used for the semantic segmentation stage, which is based on Bayesian SegNet[34]. SegNet, the base architecture of Bayesian SegNet, is a state-of-the-art network for semantic segmentation [36]. Our prior work demonstrated that SegNet performs well for hazard detection on planetary surfaces [8]. The base structure of SegNet is a convolutional encoder-decoder, which successively downsamples the input until a bottleneck layer, at which point successive upsampling stages are used to output the desired resolution. Each convolutional layer includes convolution, batch normalization, and a rectified linear unit (ReLU) activation function. Convolutional blocks



Fig. 1 Proposed approach for uncertainty-aware learning-based landing site detection.



Fig. 2 A schematic of the network architecture for semantic segmentation, which is based on Bayesian SegNet [34].

are followed by pooling/upsampling layers to achieve the encoder-decoder architecture. The network output is given as input to a softmax layer to produce the probability of each pixel to be labeled as safe or unsafe. The network is trained with cross-entropy loss.

For MC-dropout, network dropout layers are activated for both training and testing. During testing, each realization of the noisy DEM is input to the trained network *M* times, where *M* is the number of MC samples. The final prediction is given by the mean over the softmax output across the *M* stochastic samples. Formally, let \mathcal{D}_{train} be the training set, $C = \{\text{Safe, Unsafe}\}$ be the set of pixel classes, *y* be the output class associated with one DEM pixel $\mathbf{x} = (u, v)$, and $p(y = c | \mathbf{x}, \hat{w}_m)$ be the predicted probability that the DEM pixel $\mathbf{x} = (u, v)$ has the output label $c \in C$ given the network weight realization \hat{w}_m at the *m*th stochastic forward pass under random dropout. Then, the final prediction is

$$\hat{p}(y=c|\mathbf{x},\mathcal{D}_{train}) = \frac{1}{M} \sum_{m=1}^{M} p(y=c|\mathbf{x},\hat{w}_m).$$
(1)

To output the final segmentation map, we take the label *c* that corresponds to the highest $\hat{p}(y = c | \mathbf{x}, \mathcal{D}_{train})$ for each pixel \mathbf{x} .

B. Uncertainty Estimation

The uncertainty map can be estimated during testing. There are several methods for computing the uncertainty map given the network output of the *M* stochastic samples to capture aleatoric and/or epistemic uncertainty. These methods include computing the variance of the output samples, predictive entropy of the mean softmax probability, or mutual information between the predictive entropy and the posterior over model parameters [13]. Mukhoti and Gal [13] notes that predictive entropy models both epistemic and aleatoric uncertainty, which are both important for safe planetary landing. Thus, for this work, we use predictive entropy to compute uncertainty maps. The predictive entropy, $\hat{\mathbb{H}}[y|\mathbf{x}, \mathcal{D}_{train}]$, is given by [13]:

$$\hat{\mathbb{H}}[y|\mathbf{x}, \mathcal{D}_{train}] = -\sum_{c \in C} \left[\left(\frac{1}{M} \sum_{m=1}^{M} p(y=c|\mathbf{x}, \hat{w}_m) \right) \log \left(\frac{1}{M} \sum_{m=1}^{M} p(y=c|\mathbf{x}, \hat{w}_m) \right) \right]$$
(2)

Figure 3 shows an example of the outputs.



Fig. 3 Example of safety map prediction and uncertainty estimation.

From left to right, Figure 3 shows an example input noisy DEM, the ground truth label produced by running a replicated version of the ALHAT algorithm on an input true DEM without noise, the output network prediction, and the output uncertainty map. The network prediction shows the network output labeling each pixel as safe or unsafe,

reported in blue and yellow, respectively. Note that border pixels are considered invalid as the safety assessment of each pixel requires consideration of neighboring pixels, with the size of the neighborhood dependent on the lander size. These pixels are ignored during training and evaluation. The uncertainty map is obtained by calculating predictive entropy. Low uncertainty, reported in black, indicates that the network prediction can be trusted. High uncertainty, reported in yellow, indicates that the network output may be inaccurate.

C. Uncertainty-aware Safety Map Generation

It is necessary to select an uncertainty threshold, \mathbb{H}_t , for determining whether the uncertainty value should be interpreted as certain or uncertain [13]. Predictions determined to be uncertain should not be trusted. Formally, given an uncertainty threshold, \mathbb{H}_t , and uncertainty value, $\hat{\mathbb{H}}$, we interpret the prediction as uncertain and make it invalid if $\hat{\mathbb{H}} \geq \mathbb{H}_t$. Otherwise, the predicted labels are retained. Figure 4 provides an illustrative example to demonstrate the effect of varying the uncertainty threshold to determine whether predictions for each pixel are certain or uncertain. Selecting a conservative threshold leads to many accurate pixels being labeled uncertain, potentially discarding accurate predictions. With a less conservative threshold, more predictions pass through as certain and accurate. However, some predictions appear as certain although they are inaccurate, which could be dangerous for a safety-critical application. Selecting the proper uncertainty threshold is a critical consideration for uncertainty-aware safety prediction.



Key for Uncertainty Performance Image: Green – Accurate/Certain Blue – Accurate/Uncertain Light Blue – Inaccurate/Uncertain Red – Inaccurate/Certain

Fig. 4 Effect of varying the uncertainty threshold to determine if the prediction is certain or uncertain.

The uncertainty threshold is selected based on a statistic from the training set or, if the training set is too large, the validation set. Note that the uncertainty threshold should not be fully dependent on the information at the observation. For example, consider the case where the threshold is the mean value of the obtained uncertainty map. Then, within the predicted safety map, there always (incorrectly) exist some pixels interpreted certain, no matter how large the minimum value of the obtained uncertainty map is. In this study, the uncertainty threshold is set to be the mean uncertainty value across the validation set, as suggested by Mukhoti and Gal [13] for general semantic segmentation tasks. Moreover, we also studied the effect of varying the uncertainty threshold for our application in Section IV.C.3. Figure 5 summarizes the pipeline of integrating the uncertainty map and network prediction.



Fig. 5 Pipeline for integrating uncertainty map and network prediction to enable uncertainty-aware safety map generation.

IV. Experiments & Results

A. Dataset Preparation

Network training and testing is achieved using simulated terrain data for input DEMs and a replicated version of the ALHAT's hazard detection (HD) algorithm for output safety maps.

1. Simulated terrain data:

Improving from our previous work [37], this work leverages an open-source sensor simulation toolbox, BlenSor, to simulate data collection of realistic terrain with LiDAR [38]. BlenSor is built upon Blender, an open-source 3D simulation engine. Figure 6 illustrates the pipeline for creating the simulation environment and generating the simulated dataset. To create the simulation environment, we first load a Mars High Resolution Imaging Science Experiment (HiRISE) digital terrain model (DTM) into BlenSor using the Blender HiRISE plug-in [39]. The Mars HiRISE DTMs are derived from stereo image pairs collected by the HiRISE camera on the Mars Reconnaissance Orbiter (MRO) [39]. This provides a realistic base terrain collected during a real mission for generating simulated data. The original Mars HiRISE DTM is selected for the candidate future landing site near Sabrina Vallis with 1 m/pixel resolution. Due to memory limitations, we load the DTM at 25% of the full resolution. Next, we configure a time-of-flight sensor to simulate the parameters of a real landing scenario. Table 1 provides the sensor parameters of the simulated sensor. These parameters are selected to obtain DEMs with similar dimensions and characteristics of DEMs obtained through real and simulated experiments with flash LiDAR surveys for imaging planetary surfaces [20, 40–42]. For the simulation, the sensor is moved across the terrain at a fixed altitude of 500 m to collect incremental scans of the terrain. The

output from BlenSor is a point cloud. The point cloud is converted to a DEM with a resolution of 1 m/pixel using bilinear interpolation. The final DEM is cropped to a resolution of 100 x 100. BlenSor allows for the output of the true (noise-free) sensor scan, or a noisy scan. We generated a low-noise dataset with a 1- σ noise level of 1.67 cm (S-167), a moderate-noise dataset with a 1- σ noise level of 3 cm (S-300), and a high-noise dataset with a 1- σ noise level of 7 cm (S-700). Each dataset uses the same terrain and the same sensor trajectory to ensure that the main difference in the datasets is the noise level. Figure 7 shows the simulation environment during the data generation process, and Fig. 8 shows the sample output terrain obtained through our simulation pipeline.

Note that we can add any other noise models like distortion or damaged elements of the LiDAR detector for the training data, given the mission-specific observation model. The altitudes and the angles of the sensor boresight axis can also be varied so that the training data distribution aligns well with a specific situation during descent. In this paper, however, we assume that we can perform the appropriate resampling of the raw DEM from the LiDAR point cloud data (PCD). This process up-samples or down-samples the raw DEM from the PCD to have a uniform resolution that is identical to the training DEMs. The detailed resampling-based approach and its application in a simulated Mars landing scenario can be found in Ref. [43].



Fig. 6 Flowchart illustrating the pipeline for creating a realistic virtual environment and generating simulated data.

 Table 1
 Sensor parameters for generating simulated data in BlenSor [38].

Parameter	Value
Detector size	128x128
Horizontal field-of-view	12°
Vertical field-of-view	12°
Focal length	25 cm





Fig. 7 Simulation environment for data generation

Fig. 8 Example of generated terrain

2. Label generation:

To prepare ground truth labels of safety maps, we replicated ALHAT's HD algorithm [18]. ALHAT's HD algorithm calculates the maximum slope and roughness the lander would experience at touchdown by explicitly evaluating the lander geometry and the landing pad contacts to the given terrain. We replicated deterministic slope evaluation, deterministic roughness evaluation, and probabilistic roughness evaluation. For label generation, we applied deterministic slope and roughness evaluation to the noise-free DEMs. For ALHAT performance evaluation as a baseline, we applied deterministic slope evaluation and probabilistic roughness evaluation as proposed by [18], resulting in the algorithm that returns pixel-wise safety probability. The baseline performance is measured by adopting a safety threshold of 50% to assign a label of safe or unsafe to each pixel. Note that this algorithm requires evaluation of all the surrounding pixels for every aiming point, so it may not be used online for high-resolution DEM inputs without approximation.

B. Implementation Details

We generated simulated datasets of 1000 DEMs. From this data, 800 DEMs are used for training, 100 DEMs are used for validation, and 100 DEMs are used for testing, selected through random selection. Each training set contains a noised DEM and a label generated with the replicated ALHAT algorithm from the noise-free DEM. The input data is 100x100 resolution at 1 meter per pixel. The data is upsampled with bilinear interpolation to 512x512 before being input to the network. The data is also normalized between 0 and 1 based on the minimum and maximum values across the training set.

For training, we use a batch size of 8. The dropout rate for dropout layers is set to 0.5 following Reference 34. The

learning rate is 0.0001 and the momentum is 0.9. The network is trained for 10^4 epochs, which means that the full training set is seen 10^4 times. This training takes approximately 2 days on an NVIDIA GeForce RTX 2070 SUPER GPU. For testing, we use M = 8 samples for MC-dropout.

C. Results & Discussion

1. Qualitative Comparison with Visualized Output

Figure 9 shows qualitative results of the network prediction together with the baseline performance for a sample DEM with varying levels of noise: low noise of 1.67 cm in $1-\sigma$ (S-167), moderate-noise of 3 cm in $1-\sigma$ (S-300), and high-noise of 7 cm in $1-\sigma$ (S-700). Note that the uncertainty threshold is selected to be the mean uncertainty in the S-167 validation set for all cases. On the top, we show the true, noise-free DEM and corresponding safety label. We can observe that safe-unsafe border pixels have higher uncertainty, and their area increases as the noise level increases. As our method incorporates the uncertainty map in our final prediction, our uncertainty-aware safety map has decreased error for valid pixels, at the expense of increasing the number of invalid pixels. In other words, our method can cope with the increased sensor noise with the aid of the uncertainty map. On the other hand, the baseline prediction from the replicated ALHAT algorithm, which is a state-of-the-art method, degrades with increased sensor noise as reported in Ivanov et al. [18]. The baseline method is conservative, especially for the high noise input, where all pixels are labeled as unsafe and there are no viable safe landing sites. The learning-based approach can still return viable safe landing site candidates, even when the noise of the test sample is outside the distribution of the training set.

2. Quantitative Comparison with Performance Metrics

The following section presents quantitative metrics to evaluate the results: pixel accuracy, mean intersection over union, precision, and sensitivity. Each metric is described below.

Pixel accuracy, PA, is computed by

$$PA = \frac{\sum_{i=1}^{n_c} N_{ii}}{\sum_{i=1}^{n_c} \sum_{i=1}^{n_c} N_{ij}}$$
(3)

where $n_c = |C|$ is the number of classes, and N_{ij} is the number of pixels of class *i* predicted as class *j* [44]. We compute pixel accuracy across all of the valid pixels in the test set. Invalid pixels are the border pixels and, for the proposed uncertainty-aware method, the pixels with higher uncertainty than the threshold.

Mean intersection over union is a common metric for evaluating semantic segmentation. Intersection over union for a class c is computed by



Fig. 9 Qualitative sample results compared with baseline method for varying levels of input noise.

$$IoU(c) = \frac{N_{cc}}{\left(\sum_{i=1}^{n_c} N_{ci} + \sum_{i=1}^{n_c} N_{ic}\right) - N_{cc}}$$
(4)

where the numerator and denominator represent respectively the intersection and union of the predicted pixels and the ground truth of the class *c* pixels. The mean intersection over union (mIoU) is defined as the mean IoU across all classes, mIoU = $n_c^{-1} \sum_{c \in C} IoU(c)$. We compute mIoU with valid pixels only, as in the PA evaluation.

Lastly, we provide precision and sensitivity. Precision represents how reliable the predicted safe sites are and sensitivity represents how many safe sites are detected among the true safe sites. These values are calculated as following:

$$Precision = \frac{(True Safe)}{(True Safe) + (False Safe)}$$
(5)

Sensitivity =
$$\frac{(\text{True Safe})}{(\text{True Safe}) + (\text{False Unsafe})}$$
 (6)

where false unsafe include the safe pixels in ground truth that are classified as invalid due to their high uncertainty. These counts are accumulated across pixels for the full test set and the rates are computed at the end.

Table 2 shows the results for these metrics compared across different methods. The baseline method results are generated from our replicated version of the probabilistic ALHAT algorithm. The base network method results, denoted by "Base Net.", evaluate the initial network prediction obtained through MC-sampling with Eq. 1, which is the output safety map without incorporating uncertainty. The uncertainty-aware prediction is the result of our proposed method to integrate estimated uncertainty and the network prediction to provide a final uncertainty-aware safety map. The uncertainty threshold is selected to be the mean uncertainty in the validation set of the S-167 dataset.

In Table 2, the baseline achieves the best precision value across the various noise level inputs, but its sensitivity significantly degrades as the noise increases. In other words, the baseline is too conservative to detect safe sites by the severely noised DEMs. This result is consistent with the visualization in Fig. 9. In contrast, the base network detects the safe sites most as seen in the best sensitivity values, while the precision is lowered. The proposed method achieves the balanced precision and sensitivity values and also the best overall segmentation performance with the best pixel accuracies and mIoUs. The rate of pixels that pass the uncertainty thresholding, shown in the V/C Pix. % column, decreases as the noise level increases. However, the other performance metrics of the base network and the proposed uncertainty-aware methods do not degrade as much as the baseline method.

Table 2 Quantitative results comparison between the baseline replicated ALHAT method (Baseline), Bayesian deep learning prediction before uncertainty thresholding (Base Net.), and uncertainty-aware Bayesian deep learning prediction (Uncertainty-Aware). V/C pixels stand for valid and certain pixels and Pix. Acc. stands for pixel accuracy. Bold numbers correspond to the best results for each combination of the noise case and the metric.

Method	Train Noise 1- σ	Test Noise 1 - σ	V/C Pix.%	Pix. Acc.	mIoU	Precision	Sensitivity
Baseline (Replicated ALHAT)	-	1.67cm	-	0.9345	0.8294	0.9798	0.7506
Base Net.	1.67cm	1.67cm	_	0.8687	0.7347	0.6639	0.9488
Uncertainty-Aware	1.67cm	1.67cm	64%	0.9839*	0.9615*	0.9493	0.7182
Baseline (Replicated ALHAT)	_	3cm	-	0.8702	0.6663	0.9795	0.4849
Base Net.	1.67cm	3cm	_	0.8516	0.7095	0.6330	0.9510
Uncertainty-Aware	1.67cm	3cm	63%	0.9712*	0.9343*	0.9228	0.7375
Baseline (Replicated ALHAT)	-	7cm	-	0.7539	0.3790	0.9723	0.0044
Base Net.	1.67cm	7cm	_	0.8487	0.7032	0.6321	0.9284
Uncertainty-Aware	1.67cm	7cm	58%	0.9722*	0.9320*	0.9092	0.5934

* The metric is evaluated by ignoring any pixel that has higher uncertainty than the threshold. The V/C Pix.% shows the rate that passed the uncertainty thresholding. The uncertainty threshold is selected to be the mean uncertainty in the S-167 validation set.

3. Effect of Varying the Uncertainty Threshold

Figure 10 shows the precision and sensitivity of the proposed method with varying uncertainty thresholds for all three noise level cases. The selected uncertainty threshold is expressed in quantile in the S-167 validation set. For instance, the fifth data points from the left with the horizontal axis value of 0.5 represent when the uncertainty threshold is selected to be the median of the uncertainty values of the S-167 validation set. The dashed line denotes the mean uncertainty of the validation set, which is the threshold used in Fig. 9 and Table 2. We can observe that the lower uncertainty threshold makes more pixels classified invalid, which results in lower sensitivity and less safe sites detected. In contrast, the lower threshold makes precision higher and the predicted safe sites more reliable.

As a hazard detection algorithm for autonomous safe planetary landings, the goal is to provide as many safe sites as possible while maintaining the reliability of the prediction. Figure 10 shows we can achieve this by adjusting the uncertainty threshold accordingly. For instance, in the S-300 case, setting uncertainty threshold as the 40% quantile achieves 99.05% precision and 44.18% sensitivity and 50% quantile threshold with 97.83% precision and 55.07% sensitivity. Compared with the baseline of 97.95% precision and 48.49% sensitivity, the proposed method offers the flexibility to explore the trade-off between precision and sensitivity. In the S-700 case, the proposed method can outperform the baseline in terms of both precision and sensitivity. Setting the uncertainty threshold as the 40% quantile in the S-700 case achieves the precision of 97.80% while having the sensitivity of 27.12%, which significantly





Fig. 10 Qualitative sample results compared with baseline method for varying levels of input noise.

4. Effect of Dead Elements of LiDAR Detector

LiDAR systems might have a few completely dead elements, or a few damaged elements with large range errors. Such an incomplete observation also needs to be processed by the hazard detection algorithm without severely degrading the performance. Ideally, we would retrain the network with the damaged DEMs. However, our analysis shows that bilinear interpolation is sufficient for the proposed method to handle the incomplete observation when the missed range measurements are sparse and limited in number. Table 3 shows the performance of each method for the simulated damaged DEMs. We randomly dropped ten sparse pixels out of the 100 x 100 input DEM pixels. The dropped pixels are reconstructed via bilinear interpolation before the upsampling to the 512 x 512 pixels and processed by the network. The metrics evaluation for all cases including baseline, ignores any pixel around the dropped pixel with the lander size, for which the information is insufficient and we avoid landing. The pixel accuracy and mIoU for the proposed method in the third row of Table 3 ignores also the pixels with higher uncertainty than the threshold. The uncertainty threshold is selected to be the mean uncertainty value in the validation set of S-167, which does not contain the damaged DEMs. The results in Table 3 show that the dropped pixels in the damaged DEMs do not affect the prediction by the proposed method when the DEM is corrected via bilinear interpolation and the minimum neighborhood of those pixels by the lander size is ignored.

V. Conclusion

This paper proposed a method for uncertainty-aware deep learning for safe landing site detection on planetary surfaces. Despite the recent advances of artificial intelligence techniques, their applicability to safety-critical space missions has often been limited due to concerns regarding their outputs' reliability. In this paper, we demonstrated that incorporating network uncertainty into the final safety map can improve its performance of safe landing site detection. More importantly, we also demonstrated that safe sites can be reliably detected from DEMs with high noise, and from

Table 3 Quantitative results comparison for the restored input DEMs from the incomplete observation with ten randomly dropped pixels, between the baseline replicated ALHAT method (Baseline), Bayesian deep learning prediction before uncertainty thresholding (Base Net.), and uncertainty-aware Bayesian deep learning prediction (Uncertainty-Aware). V/C pixels stand for valid and certain pixels and Pix. Acc. stands for pixel accuracy. Bold numbers correspond to the best results for each combination of the noise case and the metric.

Method	Train Noise 1- σ	Test Noise $1-\sigma$	V/C Pix.%	Pix. Acc.	mIoU	Precision	Sensitivity
Baseline (Replicated ALHAT)	_	1.67cm	-	0.9350*	0.8299*	0.9796*	0.7511*
Base Net.	1.67cm	1.67cm	_	0.8687*	0.7308*	0.6715*	0.9122*
Uncertainty-Aware	1.67cm	1.67cm	63%	0.9828 [†]	0.9567 [†]	0.9488*	0.6517*

* The metric is evaluated by ignoring any pixel neighboring the dropped pixel with the size of the neighborhood dependent on the lander size. † The metric is evaluated by ignoring any pixel that has higher uncertainty than the threshold or is neighboring the dropped pixel with the size of the neighborhood dependent on the lander size. The V/C Pix.% shows the rate that passed the uncertainty thresholding. The uncertainty threshold is selected to be the mean uncertainty in the S-167 validation set.

DEMs where the sensor noise of the test samples is outside the distribution of the training samples. The proposed method can be an important step towards a reliable and robust machine-learning method for realistic spacecraft applications.

Acknowledgments

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

References

- Mustard, J., Adler, M., Allwood, A., Bass, D., Beaty, D., Bell, J., Brinckerhoff, W., Carr, M., Des Marais, D., Brake, B., et al., "Report of the Mars 2020 Science Definition Team," *Mars Explorer Program Analysis Group*, 2013.
- Martin, M. S., Mendeck, G. F., Brugarolas, P. B., Singh, G., Serricchio, F., Lee, S. W., Wong, E. C., and Essmiller, J. C.,
 "In-flight experience of the mars science laboratory guidance, navigation, and control system for entry, descent, and landing," *CEAS Space Journal*, Vol. 7, No. 2, 2015, pp. 119–142. https://doi.org/10.1007/s12567-015-0091-3.
- [3] Johnson, A. E., Huertas, A., Werner, R. A., and Montgomery, J. F., "Analysis of on-board hazard detection and avoidance for safe lunar landing," 2008 IEEE Aerospace Conference, IEEE, 2008, pp. 1–9. https://doi.org/10.1109/AERO.2008.4526301.
- [4] Huertas, A., Johnson, A. E., Werner, R. A., and Maddock, R. A., "Performance evaluation of hazard detection and avoidance algorithms for safe Lunar landings," 2010 IEEE Aerospace Conference, IEEE, 2010, pp. 1–20. https://doi.org/10.1109/AERO. 2010.5447022.
- [5] Rutishauser, D., Epp, C., and Robertson, E., "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. https://doi.org/10.2514/6.2012-5239.

- [6] Trawny, N., Huertas, A., Luna, M. E., Villalpando, C. Y., Martin, K., Carson, J. M., Johnson, A. E., Restrepo, C., and Roback, V. E., "Flight testing a real-time hazard detection system for safe lunar landing on the rocket-powered morpheus vehicle," *AIAA Guidance, Navigation, and Control Conference*, 2015, p. 0326. https://doi.org/10.2514/6.2015-0326.
- [7] Li, S., Jiang, X., and Tao, T., "Guidance summary and assessment of the Chang'e-3 powered descent and landing," *Journal of Spacecraft and Rockets*, Vol. 53, No. 2, 2016, pp. 258–277. https://doi.org/10.2514/1.A33208.
- [8] Tomita, K., Skinner, K., Iiyama, K., Jagatia, B., Nakagawa, T., and Ho, K., "Hazard detection algorithm for planetary landing using semantic segmentation," ASCEND 2020, 2020, p. 4150. https://doi.org/10.2514/6.2020-4150.
- [9] Moghe, R., and Zanetti, R., "A Deep learning approach to Hazard detection for Autonomous Lunar landing," *The Journal of the Astronautical Sciences*, Vol. 67, No. 4, 2020, pp. 1811–1830. https://doi.org/10.1007/s40295-020-00239-8.
- [10] Gustafsson, F. K., Danelljan, M., and Schon, T. B., "Evaluating Scalable Bayesian Deep Learning Methods for Robust Computer Vision," 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2020, pp. 1289–1298. https://doi.org/10.1109/CVPRW50498.2020.00167.
- [11] Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q., "On calibration of modern neural networks," *International conference on machine learning*, PMLR, 2017, pp. 1321–1330.
- [12] Kendall, A., and Gal, Y., "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" Advances in Neural Information Processing Systems, Vol. 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper/2017/file/ 2650d6089a6d640c5e85b2b88265dc2b-Paper.pdf.
- [13] Mukhoti, J., and Gal, Y., "Evaluating bayesian deep learning methods for semantic segmentation," arXiv preprint arXiv:1811.12709, 2018. https://doi.org/10.48550/arXiv.1811.12709.
- [14] Scalia, G., Grambow, C. A., Pernici, B., Li, Y.-P., and Green, W. H., "Evaluating Scalable Uncertainty Estimation Methods for Deep Learning-Based Molecular Property Prediction," *Journal of Chemical Information and Modeling*, 2020, pp. 2697—-2717. https://doi.org/10.1021/acs.jcim.9b00975.
- [15] Epp, C. D., and Smith, T. B., "Autonomous precision landing and hazard detection and avoidance technology (ALHAT)," 2007 IEEE Aerospace Conference, IEEE, 2007, pp. 1–7. https://doi.org/10.1109/AERO.2007.352724.
- [16] Brady, T., and Schwartz, J., "ALHAT system architecture and operational concept," 2007 IEEE Aerospace Conference, IEEE, 2007, pp. 1–13. https://doi.org/10.1109/AERO.2007.352725.
- [17] Striepe, S. A., Epp, C. D., and Robertson, E. A., "Autonomous precision landing and hazard avoidance technology (ALHAT) project status as of May 2010," *International Planetary Probe Workshop 2010 (IPPW-7)*, 2010.
- [18] Ivanov, T., Huertas, A., and Carson, J. M., "Probabilistic hazard detection for autonomous safe landing," AIAA Guidance, Navigation, and Control (GNC) Conference, 2013, p. 5019. https://doi.org/10.2514/6.2013-5019.

- [19] Epp, C., Robertson, E., and Carson, J. M., "Real-time hazard detection and avoidance demonstration for a planetary lander," *AIAA SPACE 2014 conference and exposition*, 2014, p. 4312. https://doi.org/10.2514/6.2014-4312.
- [20] Restrepo, C. I., Chen, P.-T., Sostaric, R. R., and Carson, J. M., "Next-generation nasa hazard detection system development," *AIAA Scitech 2020 Forum*, 2020, p. 0368. https://doi.org/10.2514/6.2020-0368.
- [21] Iiyama, K., Tomita, K., Jagatia, B. A., Nakagawa, T., and Ho, K., "Deep reinforcement learning for safe landing site selection with concurrent consideration of divert maneuvers," *arXiv preprint arXiv:2102.12432*, 2021. https://doi.org/10.48550/arXiv. 2102.12432.
- [22] Gaudet, B., Linares, R., and Furfaro, R., "Deep reinforcement learning for six degree-of-freedom planetary landing," *Advances in Space Research*, Vol. 65, No. 7, 2020, pp. 1723–1741. https://doi.org/10.1016/j.asr.2019.12.030.
- [23] Scorsoglio, A., Furfaro, R., Linares, R., and Gaudet, B., "Image-based deep reinforcement learning for autonomous lunar landing," *AIAA Scitech 2020 Forum*, 2020, p. 1910. https://doi.org/10.2514/6.2020-1910.
- [24] Furfaro, R., Bloise, I., Orlandelli, M., Di Lizia, P., Topputo, F., Linares, R., et al., "Deep learning for autonomous lunar landing," 2018 AAS/AIAA Astrodynamics Specialist Conference, Vol. 167, Univelt, 2018, pp. 3285–3306.
- [25] Di, K., Li, W., Yue, Z., Sun, Y., and Liu, Y., "A machine learning approach to crater detection from topographic data," *Advances in Space Research*, Vol. 54, No. 11, 2014, pp. 2419–2429. https://doi.org/10.1016/j.asr.2014.08.018.
- [26] Emami, E., Bebis, G., Nefian, A., and Fong, T., "Automatic crater detection using convex grouping and convolutional neural networks," *International symposium on visual computing*, Springer, 2015, pp. 213–224. https://doi.org/10.1007/978-3-319-27863-6_20.
- [27] Cohen, J. P., Lo, H. Z., Lu, T., and Ding, W., "Crater detection via convolutional neural networks," arXiv preprint arXiv:1601.00978, 2016. https://doi.org/10.48550/arXiv.1601.00978.
- [28] Silburt, A., Ali-Dib, M., Zhu, C., Jackson, A., Valencia, D., Kissin, Y., Tamayo, D., and Menou, K., "Lunar crater identification via deep learning," *Icarus*, Vol. 317, 2019, pp. 27–38. https://doi.org/10.1016/j.icarus.2018.06.022.
- [29] Wang, Y., and Wu, B., "Active machine learning approach for crater detection from planetary imagery and digital elevation models," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 57, No. 8, 2019, pp. 5777–5789. https://doi.org/10.1109/ TGRS.2019.2902198.
- [30] Downes, L., Steiner, T. J., and How, J. P., "Deep learning crater detection for lunar terrain relative navigation," AIAA SciTech 2020 Forum, 2020, p. 1838. https://doi.org/10.2514/6.2020-1838.
- [31] Palafox, L. F., Hamilton, C. W., Scheidt, S. P., and Alvarez, A. M., "Automated detection of geological landforms on Mars using Convolutional Neural Networks," *Computers & geosciences*, Vol. 101, 2017, pp. 48–56. https://doi.org/10.1016/j.cageo.2016. 12.015.

- [32] Pugliatti, M., Maestrini, M., Di Lizia, P., and Topputo, F., "On-board Small-Body Semantic Segmentation Based on Morphological Features with U-Net," 31st AAS/AIAA Space Flight Mechanics Meeting, 2021, pp. 1–20.
- [33] Gal, Y., "Uncertainty in Deep Learning," Ph.D. thesis, University of Cambridge, 2016.
- [34] Kendall, A., Badrinarayanan, V., and Cipolla, R., "Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding," arXiv preprint arXiv:1511.02680, 2015. https://doi.org/10.48550/arXiv.1511.02680.
- [35] Lakshminarayanan, B., Pritzel, A., and Blundell, C., "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, Vol. 30, 2017.
- [36] Badrinarayanan, V., Kendall, A., and Cipolla, R., "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 39, No. 12, 2017, pp. 2481–2495. https://doi.org/10.1109/TPAMI.2016.2644615.
- [37] Skinner, K. A., Tomita, K., and Ho, K., "Uncertainty-aware deep learning for safe landing site selection," AAS/AIAA Space Flight Mechanics Meeting 2021, 2021, pp. Paper–AAS.
- [38] Gschwandtner, M., Kwitt, R., Uhl, A., and Pree, W., "BlenSor: Blender sensor simulation toolbox," *International Symposium on Visual Computing*, Springer, 2011, pp. 199–208. https://doi.org/10.1007/978-3-642-24031-7_20.
- [39] McEwen, A. S., Eliason, E. M., Bergstrom, J. W., Bridges, N. T., Hansen, C. J., Delamere, W. A., Grant, J. A., Gulick, V. C., Herkenhoff, K. E., Keszthelyi, L., kirk, R. L., Mellon, M. T., Squyres, S. W., Thomas, N., and Weitz, C. M., "Mars reconnaissance orbiter's high resolution imaging science experiment (HiRISE)," *Journal of Geophysical Research E: Planets*, Vol. 112, 2007. https://doi.org/10.1029/2005JE002605.
- [40] Keim, J. A., Mobasser, S., Kuang, D., Cheng, Y., Ivanov, T., Johnson, A. E., Goldberg, H. R., Khanoyan, G., and Natzic, D. B., "Field test implementation to evaluate a flash LIDAR as a primary sensor for safe lunar landing," 2010 IEEE Aerospace Conference, 2010, pp. 1–14. https://doi.org/10.1109/AERO.2010.5447024.
- [41] Bulyshev, A., Vanek, M., Amzajerdian, F., Pierrottet, D., Hines, G., and Reisse, R., "A super-resolution algorithm for enhancement of flash lidar data," *Computational Imaging IX*, Vol. 7873, edited by C. A. Bouman, I. Pollak, and P. J. Wolfe, International Society for Optics and Photonics, SPIE, 2011, p. 78730F. https://doi.org/10.1117/12.876283, URL https://doi.org/10.1117/12.876283.
- [42] Amzajerdian, F., Roback, V. E., Bulyshev, A., Brewster, P. F., and Hines, G. D., "Imaging flash lidar for autonomous safe landing and spacecraft proximity operation," *AIAA SPACE 2016*, 2016, p. 5591. https://doi.org/10.2514/6.2016-5591.
- [43] Tomita, K., and Ho, K., "Adaptive Hazard Detection and Avoidance for Planetary Landing via Bayesian Semantic Segmentation," AIAA SCITECH 2022 Forum, 2022, p. 2207. https://doi.org/10.2514/6.2022-2207.
- [44] Long, J., Shelhamer, E., and Darrell, T., "Fully convolutional networks for semantic segmentation," 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 3431–3440. https://doi.org/10.1109/CVPR.2015.7298965.