

A Fully-Autonomous on-Board GNC Methodology for Small-Body Environments Based on Convolutional Neural Network (CNN) Image Processing and Model Predictive Control (MPC)s

[Pelayo Peñarroya](#)^{*}, Alfredo Escalante, Thomas Frekhaug, [Manuel Sanjurjo](#)

Posted Date: 15 October 2024

doi: 10.20944/preprints202410.1117.v1

Keywords: gnc; Autonomous systems; Machine Learning; Asteroid; Space Exploration



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

A Fully-Autonomous on-Board GNC Methodology for Small-Body Environments Based on Convolutional Neural Network (CNN) Image Processing and Model Predictive Control (MPC)s

Pelayo Peñarroya * , Alfredo Escalante, Thomas Frekhaug and Manuel Sanjurjo

Universidad Carlos III de Madrid, Avenida Universidad 30, Leganés - 28911, Madrid (Spain) 1

* Correspondence: pelayo.penarroya@deimos-space.com

Abstract: The increasing need for autonomy for space exploration missions is becoming more and more relevant in the design of missions to small bodies. The long communication latencies and sensitivity of the system to unplanned environment perturbations pose autonomous methods as a key design block for this type of mission. In this work, a fully autonomous Guidance, Navigation, and Control (GNC) methodology is introduced. This methodology relies on published Convolutional Neural Network (CNN)-based techniques for surface recognition and pose estimation and on also existing Model Predictive Control (MPC)-based techniques for the design of a trajectory to perform a soft landing on an asteroid. Combining Hazard Detection and Avoidance (HDA) with relative navigation systems, a Global Safety Map (GSM) is built on-the-fly as images are acquired. These GSMs provide the GNC system with information about feasible landing spots and populate a longitude-latitude map with safe/hazardous labels that are later processed to find an optimal landing spot based on mission requirements and a distance-from-hazard metric. The methodology is exemplified using Bennu as the body of interest and a GSM is built for an arbitrary reconnaissance orbit.

Keywords: GNC; autonomous systems; machine learning; asteroid; space exploration

1. Introduction

Autonomous GNC systems play a pivotal role in the success of space missions, particularly those aimed at exploring asteroids. As the space industry ventures into the cosmos, the complexities of navigating and controlling spacecraft in remote and dynamic environments demand innovative solutions. Traditional GNC methods, while effective in many scenarios, often face challenges when applied to the unique conditions presented by asteroids, such as irregular shapes, low gravity, and limited communication bandwidth.

In recent years, the integration of Artificial Intelligence (AI) techniques has emerged as a promising approach to address these challenges and enable fully autonomous GNC systems for asteroid missions. By harnessing the power of AI, spacecraft can adapt and respond to changing conditions, make real-time decisions, and navigate with unprecedented accuracy and efficiency. This paradigm shift from rule-based algorithms to AI-driven methodologies opens up new possibilities for exploration and scientific discovery.

For instance, in Origins, Spectral Interpretation, Resource Identification, and Security–Regolith Explorer (OSIRIS-REx), a mission undertaken by NASA and aimed at studying Bennu's composition and history [1], relative navigation techniques depend on non-autonomous landmark identification, which implies human effort and the loss of a higher degree of automation for the navigation system. In [2], a more robust technique was presented that used the shadows cast on the terrain by different features to detect them. Additionally, their method could also infer surface properties such as localization and sizing of the surface landmarks detected.

Nevertheless, OSIRIS-REx offers too a notable example of AI's impact on autonomous GNC systems. This mission, relied heavily on AI-based techniques for navigation and hazard avoidance.

The use of machine learning algorithms enabled the spacecraft to autonomously identify safe landing sites, assess surface features, and execute complex manoeuvres with precision and reliability.

On top of that, the accuracy that is typically required for landing and Touch-and-Go (TaG) phases can only be achieved with high-resolution shape models. Achieving these models is a time-consuming task that usually defines the initial phases of the mission at arrival to the body of interest. Trying to tackle this problem, [3] studied a vision-only autonomous pipeline used for terrain-relative positioning, landmark matching, and dense mapping that would grant autonomy to the mission by performing all of that on-board. Basing their work on stereophotogrammetry techniques, the pipelines do not require any *a priori* knowledge of the environment.

In practice, this pipeline ingests images generated by the spacecraft's optical sensors and feeds them to the Structure-from-Motion (SfM) algorithm that, after performing stereo matching, is able to produce a dense point cloud and a mesh for the shape model of the body. An important aspect of this sequence is that the final shape model is built as the input data is acquired, and does not need to wait for a batch of images to process them all. These sequential approaches are highly desirable when autonomy is in the discussion because they provide a much more flexible and adaptable approach that is robust to unexpected events.

In addition to OSIRIS-REx, several other missions have highlighted the effectiveness of AI in space exploration. For instance, the Mars rovers Spirit [4], Opportunity [5], and Curiosity [6] have successfully employed AI for autonomous navigation and scientific discovery on the Martian surface. Similarly, the ESA's Rosetta mission [7] utilized AI algorithms for comet rendezvous and landing, demonstrating the versatility of AI-based GNC systems across different celestial bodies.

While traditional GNC methods rely on predefined rules and models, AI-based methodologies leverage data-driven approaches to adapt and learn from the environment. The generation of synthetic data needed to train the AI-models has become a huge field of research as AI-techniques have become more important in space exploration, which is a field for which not large amounts of real-world data are available. To address this problem, researchers have spent the last years developing pipelines to procedurally generate datasets of images that can be used to train CNNs for Image Processing (IP). More general image generation pipelines, such as the one shown in [8], address the challenges of data quantity and quality in training CNN-based sensors for applications like medical analytics and autonomous driving. In [8], the authors propose a pipeline for augmenting image datasets using computer graphics and generative neural networks, which allows for the creation of realistic and varied synthetic data. This approach has improved the accuracy of CNN-based sensors compared to using only real-world data.

More related to space images, [9] introduces an open-source deep learning pipeline designed for on-orbit visual navigation, featuring two custom tools: one for generating labelled synthetic data using Blender, and another for dataset curation and model training. Similarly, [10] investigates data-driven image processing options for a CubeSat mission around a binary asteroid system, focusing on centroid methods and artificial intelligence techniques. The study introduces an enhanced center of brightness method and assesses the performance of various methods, finding that CNNs outperform others. Using a tailored pipeline for image generation and model training, this work highlights the importance of filtering techniques in extracting spatial information from images.

Machine learning algorithms, trained on vast datasets of simulated and real-world scenarios, enable spacecraft to make informed decisions in uncertain and dynamic environments. Reinforcement learning techniques further enhance autonomy by allowing spacecraft to optimize trajectories and control policies. In [11], the authors introduce the Autonomous Rendez-vous Transformer (ART), which uses Transformers to learn near-optimal policies and warm-start optimisers. ART outperforms other methods in predicting fuel-efficient trajectories and achieves efficient, constraint-satisfying solutions with fewer iterations and comparable runtime.

Autonomous relative navigation is a very hot topic in space exploration and many works based on image-processing techniques have been proposed in recent years. Works like the previously mentioned

[10] or [12], which presents BennuNet; a hybrid neural network method for estimating spacecraft relative position and attitude near minor bodies using monocular vision. BennuNet targets the asteroid Bennu with synthetic training sets from the OSIRIS-REx mission. Modern architectures and data augmentation techniques improve performance over previous CNNs, maintaining high accuracy with both synthetic and real images.

All of the previously mentioned works have in common their intention to push further the levels of mission autonomy, specifically in terms of environment understanding. Given that autonomy is of particular interest in missions that try to go as close as possible to the body they are exploring, landing scenarios pose an especially challenging topic.

Nevertheless, all of these pieces of research leave a gap in the step in between navigation and guidance: the decision-making involving the selection of the landing spot. This segment of the GNC process is called HDA and has been one of the main pillars of autonomy on board, for specific cases where a spacecraft interacts with a body or surface.

One of the first works related to HDA was introduced in [13], which aimed to develop descent and landing GNC systems to meet requirements for global lunar surface access, with landing accuracies ranging from hundreds of meters for crew missions to tens of meters for outpost missions. Later on, in works like [14] or [15], HDA techniques were pushed forward, even including AI for image processing tasks. Recently, [16] shows how passive systems, i.e. those using only images as input, are capable of performing HDA accounting for features, shadows, and slopes; using only semantic segmentation techniques, based on CNNs.

If a spacecraft could orbit around a small body and autonomously decide where to land, combining that with autonomous environment recognition and autonomous trajectory guidance and control would result in a fully autonomous GNC loop that would be able to plan a complete landing sequence making informed decisions and assessing risks and potential hazards by itself.

One of the key elements that most of the aforementioned techniques need is a shape model of the body of interest. Building such a model could take a relatively large computational effort, which could go against the autonomy of the mission. Nevertheless, this is currently being researched and works such as [17] feature a limb-based method to extract edge information from images to build low-to-medium resolution shape models from multiple viewing angles, leveraging shape-from-silhouette concepts. By testing this method on simulated images of asteroid Itokawa using Blender, the study examines its effectiveness in handling surface concavities, varied viewing angles, resolutions, and camera geometries, highlighting the dependence of model density on image selection and the necessity of multi-axis rotational analysis for accurate navigation. Following that line of research, [18] presents a novel Pole-from-Silhouette (PfS) algorithm, robust to non-zero sun phase and shadow illumination conditions, crucial for the optical navigation pipeline of autonomous spacecraft approaching unknown small bodies. The PfS algorithm estimates the rotation pole and 3D shape of a small body using illuminated silhouette pixels, rotation rate, spacecraft attitude, and relative distance. Detailed numerical simulations and sensitivity analyses demonstrate the algorithm's effectiveness across various scenarios and target bodies.

With that in mind, this paper proposes a novel methodology that combines multiple AI-based techniques to perform fully autonomous GNC around small bodies. In particular, the results shown in the next sections will focus on Bennu. By integrating machine learning, reinforcement learning, computer vision, and neural networks, the proposed approach aims to overcome the limitations of traditional methods and enable spacecraft to navigate, explore, and interact with asteroids with unprecedented autonomy and efficiency.

It is important to clarify that the primary objective of this work is not to develop or showcase the performance of individual components within the GNC system used in the simulations presented in later sections. Instead, this paper seeks to address a critical gap in the field: the integration of existing, validated AI-based models into a cohesive, autonomous, end-to-end GNC system. While significant progress has been made in developing standalone models for various aspects of GNC, there

is a lack of comprehensive frameworks that bring these components together in a fully integrated manner. This paper contributes to filling that gap by not only integrating these modules but also developing the necessary infrastructure to ensure seamless operation between components, a robust simulation environment, and the surrounding support systems required to evaluate the entire system's performance.

By focusing on this integration, the work provides a methodology for building and testing autonomous GNC systems that can effectively utilize the advancements in AI and control algorithms, offering a path forward for more practical, real-world applications.

In the following, Section 2 offers a comprehensive overview of the proposed methodology, highlighting its key components, implementation strategies, and potential impact on the future of asteroid exploration. Section 3 introduces the concept of safety maps and explains how they are generated. After that, Section 4 demonstrates the effectiveness and robustness of a fully autonomous methodology through simulations. Finally, Section 5 closes out this paper by providing final thoughts on the results obtained and adding a critical evaluation of the challenges that are yet to be tackled.

2. Methodology

As described above, one of the main weaknesses of space exploration missions is their very high latency schedules and the large uncertainties present in the scenarios they take place. Traditionally, each of these impediments required instruction or some sort of validation from Earth-based facilities, which hinders the development of these missions, making them not resilient nor robust to unforeseen events or unmodelled effects. Each of these contacts poses a bottleneck for the mission timeline and a huge toll on the risks that can be assumed in the mission in aspects such as proximity to the body of interest, trajectory corrections, or mission robustness, among others.

Ideally, to remove all of these bottlenecks, a spacecraft would need to be able to tackle all unplanned events and gather all necessary information about the environment by itself. That would make ground contacts not necessary or at least not so with the frequency that is traditionally required. Tasks such as terrain recognition, relative navigation, or trajectory planning are usually highly dependent on interaction with Earth facilities and telecommanding. However, providing the spacecraft with the technology necessary to perform all of these tasks autonomously would mean a great increase in mission execution performance. In turn, the operational complexity of the mission would benefit hugely since many of the tasks that are typically done on Earth-based facilities could be just avoided or, at the very least, see their execution frequency decreased.

However, having a spacecraft gather all this information autonomously does not solve the problem. The biggest step toward full autonomy comes after the information is collected: the spacecraft would need to make informed decisions based on the data gathered and design and plan the steps to be executed to achieve a certain goal. This element of the chain is a very challenging task. Designing an algorithm or system capable of dealing with all unforeseen events is borderline impossible to achieve with traditional techniques. Thankfully, the development of AI-based techniques in the last decades has been a great step forward towards the solution to this issue. For instance, [19] shows how applying Large Language Model (LLM)-based autonomous agents for decision-making based on user prompts to space GNC problems is feasible. The solution was developed within Kerbal Space Program Differential Games (KSPDG) challenge, which involves creating autonomous agents for satellite maneuvers in non-cooperative space operations. Utilizing prompt engineering, few-shot prompting, and fine-tuning, their LLM-based agent ranked 2nd in the competition, pioneering the use of LLMs in space research. AI-based techniques are very good at inferring information from scarce data (when properly trained) and offer more flexible and comprehensive solutions to complex or badly posed problems, especially in very uncertain problems.

For these reasons, this work presents a novel methodology that merges HDA, relative navigation, and trajectory optimisation techniques to build a full GNC system meant for small-body environments. This system is able to select a safe landing spot autonomously and design the trajectory and guidance

sequence required to land in the selected region. The whole process can be summed up in the following steps:

- Insertion into an observation orbit,
- Body’s surface image acquisition,
- Generation of a global safety map of the body,
- Landing site selection, and
- Landing trajectory design and guidance.

Space exploration missions usually have several orbital segments dedicated to learning about the body they are orbiting that lasts for weeks if not months. These trajectories entail complicated geometries that are highly optimised and tailored to mission needs. For instance, OSIRIS-REx [1] included several orbital phases at different ranges that are optimised for the scientific and operational objectives of the mission.

For the sake of this work, the reconnaissance trajectory design is simplified, since it is not the main focus of this paper. A polar orbit is selected to make sure that all latitudes and longitudes of the body can be observed. The orbital elements are shown in Table 1 and a depiction of the trajectory followed (in body-fixed coordinates) is shown in Figure 1.

Table 1. Orbital elements for the reconnaissance trajectory around Bennu.

Parameter	Value
Semi-major axis [km]	566.13 (2 Bennu radii)
Eccentricity [-]	0.0
Inclination [°]	110.0
RAAN [°]	0.0
Argument of Perigee [°]	0.0
True Anomaly [°]	270.0
Initial Epoch (UTC) [-]	2022-Jan-01 00:00h




Figure 1. Trajectory followed by the spacecraft during the reconnaissance orbit chosen. The square represents the beginning of the trajectory, and the circle its end. Coordinates represent body-fixed position and the shape model used for Bennu is plotted in the centre.

From such orbit, images are generated using a pipeline based on pyrender. A sampling rate of 2000 seconds is chosen, and the images generated are immediately queued for processing. This processing can be divided into two parallel lines.

On one hand, the image is used as an input to the relative navigation module, based on BennuNet, which was introduced in the previous section [12]. BennuNet is a set of CNNs capable of estimating the pose of a camera using monocular images. By training with synthetic images generated from multiple geometries, covering the whole range of camera positions, attitude, illumination conditions and camera field-of-views, the CNNs can learn the non-linear transformation from the 2D space (for a grayscale image) to the 6D pose space (three coordinates for the Cartesian position vector and three Euler angles with respect to the target body fixed frame). Compared to traditional feature or landmark-based algorithms for optical relative navigation, the CNNs offer a lightweight solution that can be uplinked to the spacecraft and is robust to adverse illumination conditions, maintaining accuracy for low and high extremes of the illumination phase angle.




Figure 2. Depiction of camera pose estimation with respect to target body fixed frame.

The set of CNNs is organised in two levels. First, a high-level classification network estimates the sector of the discretized 3D space around the target based on the input image. After testing with multiple levels of discretization, 4 divisions latitude- and 8 divisions longitude-wise, resulting in 32 sectors yield overall best accuracy. Based on the estimated sector, the image is then provided as input to one among a set of low-level regression networks, each of them trained with images covering only one sector of the discretized space (with overlapping between them to adapt for missed predictions at the sector boundaries). This low-level regression CNN provides as output the spacecraft pose vector. For this work, the attitude is assumed to be known from other sources, therefore the output of the CNN is only the three coordinates of the relative position vector with respect to the target body fixed frame.

On the other hand, hazards are extracted from the image using astroHda, which is a tool developed in [16]. AstroHda aims to forecast hazards for any small celestial body, leveraging the broad generalization abilities inherent in deep learning techniques such as CNNs. It features an architecture based on Residual Network (ResNet)-34 [20] and it takes advantage of learning techniques such as transfer learning [21]. The dataset used to train the model was generated using a pipeline that produces images using Blender.

AstroHda uses semantic segmentation to detect hazards in three layers, namely shadow detection, feature detection, and slope estimation. Figure 3 shows an example of an input image, the three layers used, and the final safety map obtained. It is designed as a passive HDA system, i.e., it only takes observations from visual instruments, such as cameras. Freeing the system from active instruments such as LiDARs means a considerable reduction in weight and cost for the spacecraft while making the system much more accessible to small satellites.

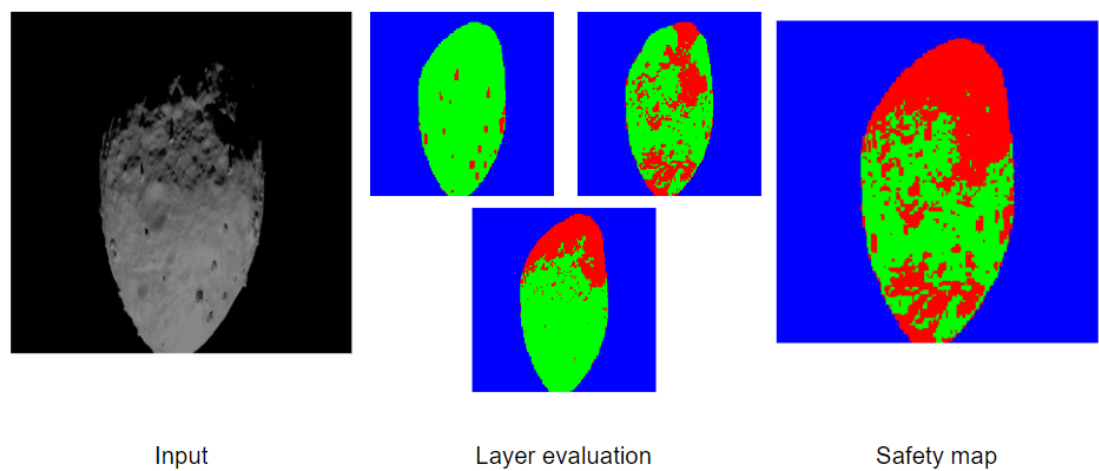


Figure 3. On the left-hand side, the original image. In the centre the three layers analysed: top-left represents feature detection, top-right represents slope estimation, and bottom represents shadow detection. On the right-hand side the safety map obtained from composing the three layers.

At this point, two post-processed inputs are available: the pose estimation for the spacecraft, and the safety map corresponding to the section of the body observed in the image. To project the hazard information onto a flat projection of Bennu’s surface, the light rays that the camera is taking need to be projected onto the surface of the mesh representing the shape model of Bennu, by means of ray-tracing techniques. In practice, this process translates pixel coordinates to longitude-latitude coordinates. Figure 4 shows the basic elements involved in ray tracing techniques. There, it can be observed how the red lines coming from the camera and pointing at a particular point on the surface of the scene object are projected onto the image plane. Section 3 describes this process in detail.




Figure 4. Ray-tracing basic elements.

By iterating the previous steps and projecting more and more safety maps onto the body’s surface, a GSM is obtained. The resolution of the GSM can be tailored to the needs of the mission, in particular

to the landing spot selection accuracy needed and to the size of the images taken and the safety maps generated by astroHda.

The next step is to choose a spot to land on. To do that, selection criteria need to be defined. One of the most straightforward options is to simply select the safe point that is surrounded by the largest amount of safe points, i.e., the safe coordinate that is furthest away from any non-safe point on the surface of the body. Other criteria might also include scientific requirements or mission-specific needs. Overlapping several criteria in a filter-like fashion could lead to more constraining conditions, which might end up having no feasible options for landing selection, so care must be taken if a landing is to be performed.

With this approach, partial evaluations can be performed while the GSM is completed. Of course, the more complete the map is, the better informed the final spot selection will be.

Once the landing spot has been selected, a descent to the desired target may be initiated. However, model uncertainties in these environments are significant due to the inherent challenges of modelling small body dynamics. Additionally, the models need to be computationally simple to the extent that it is compatible with space-graded computational resources. Similarly, without the presence of an established positional reference system, it is challenging to drive the navigation covariance towards zero. In order to achieve a safe landing, it is essential that the GNC system is robust towards these modelling and navigation uncertainties.

MPC is one promising approach for autonomous GNC in Aerospace [22]. It has inherent replanning abilities, inclusion of state and thrust-based constraints, and is optimal with respect to a cost function. Historically, MPCs have not been used in safety-critical scenarios where uncertainties are a significant factor, as uncertainty can drive MPC-based systems to become unstable. Several alternate MPC formulations have been developed over the years that aim to improve MPC in environments with high uncertainty, with Stochastic MPC and Tube-based MPC being the two main directions. Stochastic MPC aims to minimize the probability of constraint violations, given an uncertainty distribution of the model disturbances. With this, stochastic MPC can ensure probabilistic constraint satisfaction and convergence. Tube MPC is a more conservative approach that aims to bound all possible evolutions of the state and disturbances. As opposed to the probabilistic definitions of uncertainties in stochastic MPC, the uncertainties are described as an uncertainty variable bounded to a closed, convex set. With the design of an auxiliary controller, it is possible to ensure that the evolution of the state is bounded to an invariant set centred on an MPC that assumes an undisturbed system. With the simple assumption that the model and state estimation uncertainties are bounded, Tube MPC can guarantee the convergence and stability of a given system. These abilities make Tube MPC particularly attractive in safety-critical environments that include uncertainties.

Two separate phases are considered for a safe descent. An initial circumnavigation such that the spacecraft is positioned within a convex descent cone with its apex at the desired landing area. In this circumnavigation phase, an ellipsoidal Keep-Out Zone (KOZ) is enforced to ensure that the spacecraft does not collide with Bennu. The tube MPC has particular advantages that enable the linearisation of the non-convex KOZ. With this, only a simple Linear Quadratic optimization Problem (LQP) is required to be solved at each control iteration.

Once sufficiently within the descent cone, the GC system switches to the descent phase. Here, the descent cone may be constructed to include dangerous terrain or areas that have been mapped out previously. For the descent, the linear model used for the MPC is a linearisation of the estimated dynamics at the desired landing spot. With this, the uncertainties are at their smallest where most precision is needed. Additionally, no linearisation of the state space is required, as the descent cone has been constructed to be convex. With this, the tube MPC for descent is also reduced to a similar LQP as in the circumnavigation phase.

With well-constructed outer bounds of both the navigational and modelling uncertainties, the tube-based MPC can ensure safe circumnavigation and descent to any desirable landing area from any initial position.




Figure 5. Example of nine separate runs starting on a 3x3 grid around Bennu. All targeting the same initial position [23].

If all these elements are chained together, it is possible to perform autonomous navigation, generate a GSM of the body of interest, make an informed decision about where to land, design the trajectory required and successfully complete the landing sequence. From these tasks, the generation of the global safety maps and the development of the infrastructure to link all the elements in the chain are they key contribution of this work. Thus, the next section will describe more in detail the technical specifics of the different elements developed in this paper.

3. Global Safety Map Generation

The generation of GSMs is the process of adding positioning information to a safety map obtained from astroHda. Safety maps lack any kind of information about the position or orientation a certain image is taken from. This is a design decision intended to make maps more versatile and adaptable to multiple scenarios with differently shaped small celestial bodies. In this paper, maps from astroHda [16] are coupled with pose estimations from BennuNet [12], making it possible to understand which section of Bennu is being processed by astroHda.

The GSM is defined as a flat projection of longitude and latitude coordinates at the surface of the body, so the process described in this section aims to transform an image of a 3D object in space into a set of observed patches of land on the surface of Bennu. The first step to achieve that is to understand where the camera is pointing at in terms of local surface coordinates.

This is done using SPICE's ray tracing functionalities. Provided a position and attitude of the camera in the body-fixed reference frame a set of rays can be generated that go through the pixels on the safety map generated by astroHda, as it was depicted in Figure 4.

The resolution of astroHda's map, thus, determines the amount of rays that need to be traced. If these vectors intersect with the shape model, the longitude and latitude of the intersection between

the body's surface and the pixel ray is computed and the hazard label in the safety map is stored on that combination of longitude-latitude. Doing this repetitively for a sequence of images would ideally fill out a global map for the surface of Bennu.

The resolution of the GSM is set independently of the resolution of the input images and can be adjusted to the level of accuracy required by the mission. Higher resolutions will translate into longer reconnaissance stages but will increase the certainty about the hazards on the surface of the body.

During the addition of images, the same patch of land may be observed more than once. If, for instance, one of these observations happens in poor illumination conditions, there might be two different hazard labels associated with the same lat-lon element in the GSM. To solve this, a priority scheme needs to be imposed.

There are two different approaches to tackle this issue: rely on astroHda or stay on the conservative side. To minimise the risks, one would go for the latter and label as "hazardous" any pixel that has been detected as a "hazard" at least once. This makes the GSM obtained more robust and makes sure that no pixel is a false-positive (positive intended as "safe") detection. However, reconnaissance orbits usually try to observe the body from multiple angles and locations around the small celestial body and it is quite common that some of the images taken will happen to be from the dark side of the body, resulting in images with very poor illumination conditions. In the worst case, if an image is taken at the moment when the observed body is eclipsing the Sun, the safety map will result in an almost fully hazardous map, ruining many observations taken from better illumination conditions. Figure 6 shows how different illuminations can affect the generated image, even if it is taken from the exact same position with respect to Bennu.




Figure 6. Images generated from the same body-fixed coordinates around asteroid Bennu, with different illumination conditions, i.e., sunlight coming from different positions.

A middle-ground solution for this would be to include the geometry of the Sun in the process. By doing so, and since the information is available thanks to BennuNet, images taken from poor Sun-facing angles can be discarded. While this option would solve some of the problematic situations, it still would not resolve all the possible troubling cases. For instance, an image taken from a favourable Sun-face angle could still show the shadow cast by a feature on the surface of the body. Pixels inside that shadow would still be labelled as "hazardous", and would remain so even if future images would say otherwise.

With that in mind, the option chosen for this method is to trust the networks in charge of producing the safety maps, i.e., if a patch of land is detected as "safe" at least once will remain as safe for the whole map completion process. This hierarchical alternative is less robust when it comes to false-positive avoidance but offers some design advantages such as detaching the process from the modules it uses and a more generalist behaviour that needs no position-dependent pre-processing of the information.

4. Simulation Scenario and Results

The scenario used for the simulations features asteroid Bennu as the observed body. The selection of Bennu is motivated by the use of BennuNet, but the methodology would remain intact if any other small celestial body is to be studied. The model used for Bennu is shown in Figure 7.




Figure 7. Shape model of asteroid Bennu used for the simulations. The reference radius is 283.065 m.

The simulation will replicate the orbital conditions of a mission in the observation phase, i.e., in the mission segment in which the target body is circumnavigated and images of it are taken to learn about its surface properties and features. For the sake of simplicity, since orbit design is not the subject of this work, a near-polar orbit is chosen. Near-polar orbits offer the benefit of being able to scan the surface of the body, as long as they do not coincide with a particular set of frozen orbital elements that would make the orbit fly over the same ground track continuously. The orbital elements and epoch chosen are summarised in Table 1 and the trajectory is plotted in Figure 1. Notice how the Right Ascension of the Ascending Node (RAAN) has not been adjusted to match any particular stable illumination conditions, like the ones obtained in typical terminator orbits. This is intentionally done so that the images taken also include situations in which the illumination conditions are poor to analyse how those image acquisitions affect the evolution of the GSM.

Other than that, light-source direction comes from the position of the Sun with respect to Bennu at each epoch, and the dynamic model used for the orbit propagation is Keplerian, both to speed the simulation up and because orbital dynamics do not affect the validity of the methodology.

The image acquisition rate has been set to 2000 seconds and the duration of the propagation to 24 hours. This combination results in a total of 43 images. To exemplify the type of images obtained, a sample of the images obtained can be observed in Figure 8, going chronologically from left to right and from top to bottom, for a total of 24 images out of the 43 generated.




Figure 8. Series of images acquired from the spacecraft when flying along the reconnaissance orbit described in Table 1. Epoch IDs are numerically ordered and separated by a 2000 s timestep between consecutive epochs.

The images generated include a variety of different illumination conditions and correspond to an almost complete coverage of the surface of Bennu. As it was explained in the previous section, these images are processed by BennuNet and astroHda to estimate the area of the body that is being observed and the hazards present upon it, respectively. To depict how this process takes place and the evolution of the GSM, the images chosen to be depicted in Figure 8 correspond to a particular time series of images that are particularly helpful in showing the behaviour of the presented methodology. The captions for the images provide information about the epoch in which the image was acquired, with timesteps of 2000 seconds.

As introduced in Section 3, the images shown in Figure 8 are passed on to astroHda and BennuNet to obtain their HDA map and relative position estimation, respectively. Figure 9 shows the generated safety maps for each of the images acquired (see Figure 8). These maps account for illumination conditions, slopes, and features detected on Bennu and flag them as safe (in green), unsafe (in red), or deep-space (in blue). A more thorough description of the labelling process is given in [16].

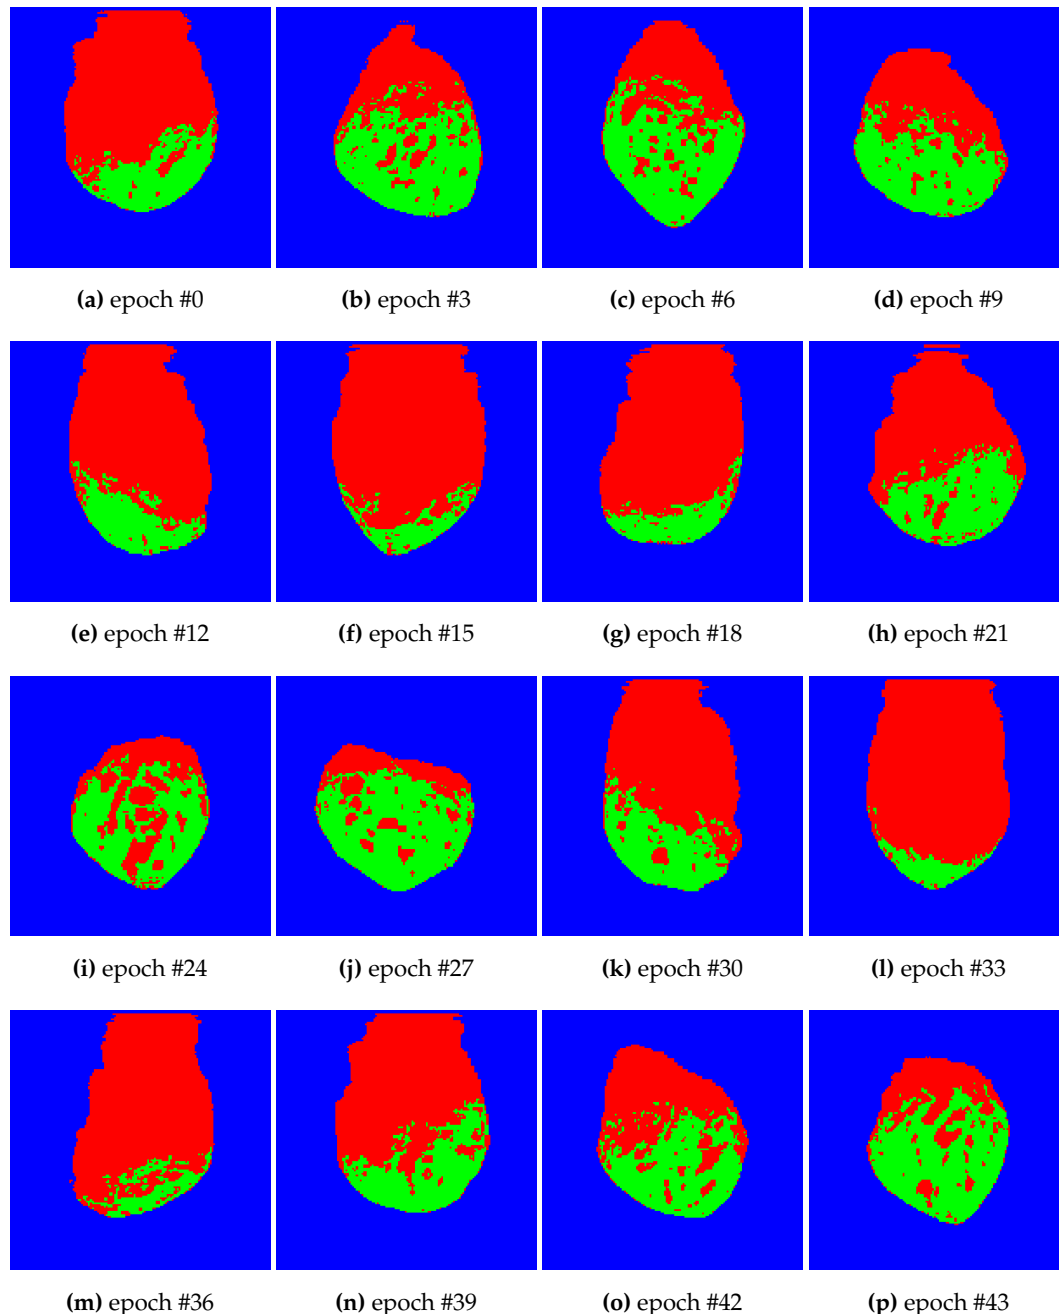


Figure 9. Series of safety maps generated from the images acquired and shown in Figure 8. Epoch IDs are numerically ordered and separated by a 2000 s timestep between consecutive epochs.

The maps shown in Figure 9 are then projected into the GSM of the body sequentially, i.e., each image acquired is processed by both astroHda and BennuNet and, following the ray-tracing techniques described in Section 2, the safety maps are projected onto a 2D map of Bennu, whose axes are its longitude and latitude. Figure 10 shows the evolution of such a map for the three initial images acquired. There, it can be observed how the pixels of the safety map are projected onto a map of

resolution 1.0°, populating the GSM with information about surface hazards. Following the priority scheme described before, any point that is observed once and deemed safe will remain so for the rest of the evolution.

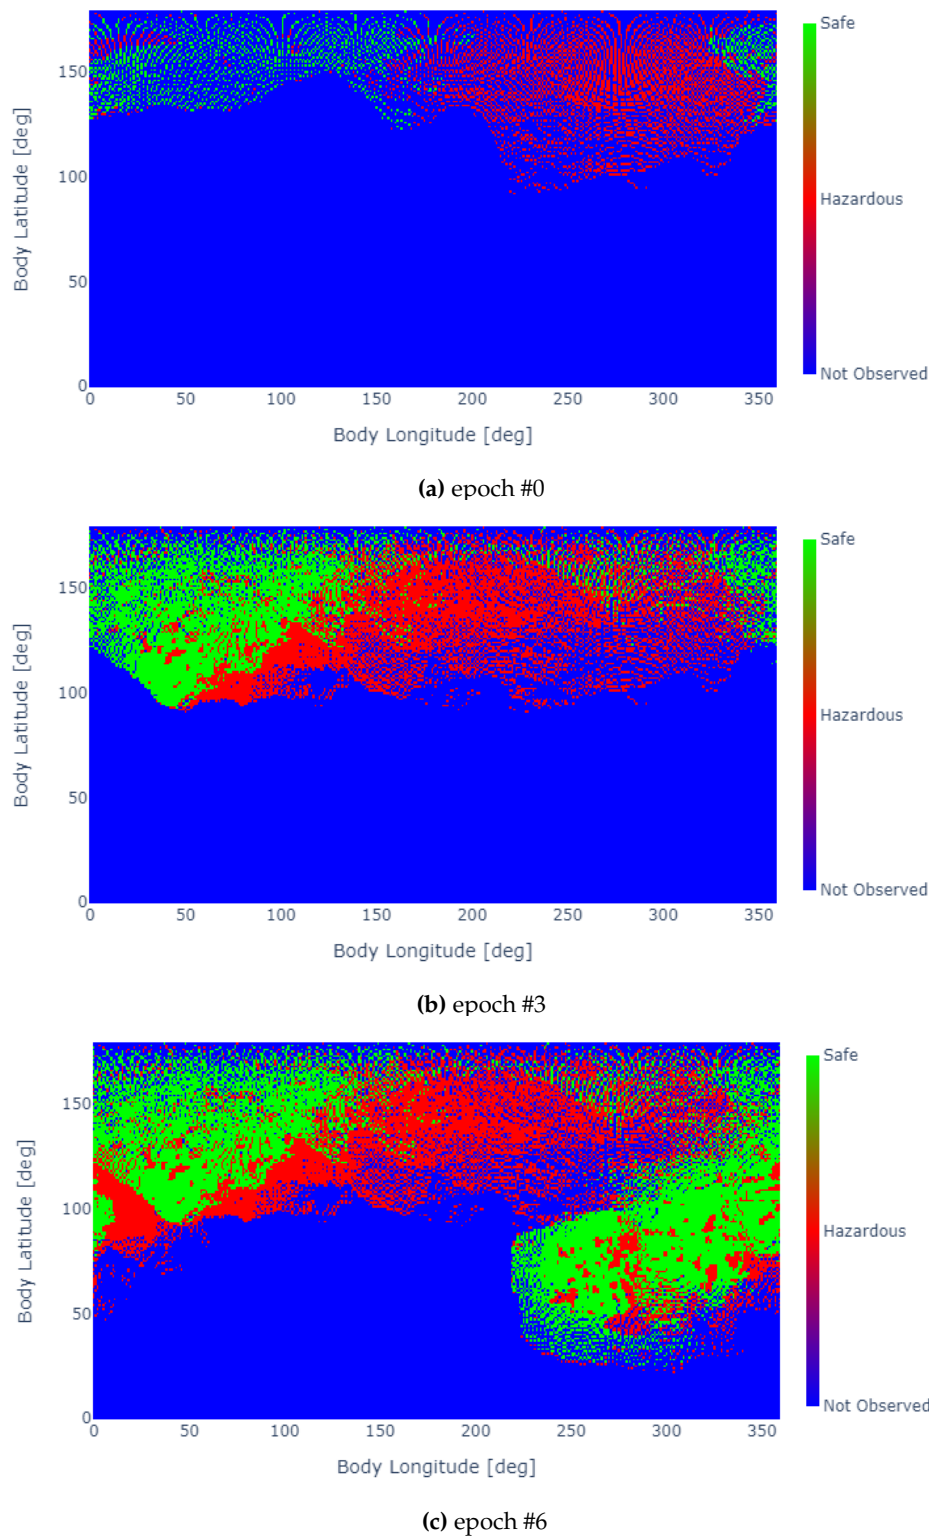



Figure 10. GSM evolution as images acquired are processed and projected onto the surface of Bennu. Epoch IDs are numerically ordered and separated by a 2000 s timestep between consecutive epochs (Initial three epochs shown).


Progressively, the GSM starts to fill up and, if the surveillance orbit is capable of observing the whole surface of the body, all longitudes and latitudes are assigned a safety value that could be:

- blue: deep-space,
- red: hazard, or
- green: safe.


Figure 11 shows a more intermediate evolution of the GSM. There, it can be observed how by epoch 18th, most of the surface of the body has been observed, although a large area on the upper latitudes of the body has only been observed under poor illumination conditions. As the surveillance orbit continues to progress, more information becomes available and even that gap is filled. Some spots along the equatorial latitudes start to take a more fixed shape: they represent features on the surface that will be flagged as hazards.



(a) epoch #6



(b) epoch #18



(c) epoch #30

Figure 11. GSM evolution as images acquired are processed and projected onto the surface of Bennu. Epoch IDs are numerically ordered and separated by a 2000 s timestep between consecutive epochs (three intermediate epochs shown).

At the end of the simulation, the GSM shown in Figure 12 is obtained. Most of the surface has been observed by that point and the main hazard areas are identified by the algorithm. Very extreme latitudes still have unobserved or hazardous labels, due to the inclination of the orbit and the fact that the body's rotational axis is quite perpendicular to the ecliptic, making it challenging to observe those

polar regions with adequate illumination conditions. Of course, this process can last indefinitely, until observations fully cover the map, but the progression of the map becomes less beneficial as fewer gaps remain.




Figure 12. GSM at the end of the last epoch of the simulation.

Figure 12 shows the final map obtained after 44 images were acquired during a propagation time of 24 hours. Once that map is obtained, a landing spot needs to be chosen. There are many possible criteria that can be used to select a certain landing spot and their motivation could arise from scientific, operational, or feasibility reasons, among others. All the different criteria to be considered can be easily overlapped to filter out areas where landing is not an option, and then select the optimal spot within the filtered subspace. Since it is not the purpose of this work to represent a specific use-case of this algorithm, no filters are applied and the engine of the landing spot selector is exposed.

This engine is based on hazard proximity, and evaluates, for each map pixel, the minimum distance to a hazardous area. Thus, Figure 12 becomes the heatmap shown in Figure 13. There, it can be observed how the areas where hazards are detected are shaded darker since their distance to a hazard would be zero. There are four regions where the distance-to-hazard metric is higher, namely:

- Long 40° - Lat 70°
- Long 40° - Lat 110°
- Long 200° - Lat 70°
- Long 340° - Lat 100°




Figure 13. Heatmap for the final GSM representing the minimum distance to a hazard from each pixel.

From all of these, the algorithm shows that the best spot corresponds to a landing spot at 36° longitude and 117° latitude, with a distance-to-hazard of 12 cells. It must be taken into account that the resolution of the map will affect the absolute magnitude of the distance-to-hazard metric, but not the qualitative results or the selection of the landing spot.

Once the landing site is selected, its coordinates are transferred to the trajectory optimization block, which employs the MPC-algorithm described previously to compute the optimal landing trajectory. The MPC is responsible for ensuring that the landing is both smooth and accurate by continuously adjusting the trajectory based on the spacecraft's current position and velocity. As introduced in Section 2, the MPC first identifies a point directly above the selected landing site, which serves as an intermediate target. This target point, also referred to as a hover point, is used to simplify the trajectory planning by breaking the descent into distinct phases. The spacecraft initially navigates towards this point before commencing its final descent to the surface. This fictitious waypoint is crucial for managing the spacecraft's speed and orientation, allowing for precise control during the critical landing phase.

Figure 14 provides a visual representation of this process, showing how the trajectory, starting from the darker blue end of the line, first guides the spacecraft to the hover point, until it enters the region defined by the pyramid depicted in blue in Figure 14. From this elevated position above the landing site, the MPC recalculates the descent path, ensuring that the final approach is safe and within acceptable landing tolerances. This method allows the system to correct for any discrepancies between the planned and actual flight paths, enabling real-time adjustments that improve landing accuracy and safety.




Figure 14. Landing trajectory towards the selected landing spot, time represented by the gradient from dark blue (initial state) to red (landing spot). The sphere around Bennu shows the region protected by Bennu's reference radius. The pyramid used to define the different optimisation phases is also shown in dark blue [23].

Once the spacecraft enters the predefined region within the optimization pyramid, the final position of the trajectory is set to the selected landing site. At this point, the MPC algorithm shifts its focus from intermediate waypoints to guiding the spacecraft directly towards the landing spot. This transition marks the beginning of the final descent phase, during which the spacecraft's trajectory is fine-tuned to ensure a smooth and controlled approach. The optimization pyramid, which represents the feasible space for landing, plays a crucial role in this phase by constraining the trajectory within safe parameters as the spacecraft approaches the surface.

This process is depicted in more detail in Figure 15, where the sharp turn towards the final landing site, marked in red, becomes clearly visible. The turn signifies the point at which the spacecraft pivots from the hovering phase to the final descent. During this stage, the MPC algorithm continuously monitors and adjusts the trajectory, ensuring that the spacecraft maintains its optimal path and speed. Any deviations from the planned trajectory are corrected in real-time to guarantee a precise and safe landing. This approach enables a high degree of accuracy in the final moments of the descent, reducing the likelihood of landing errors and improving mission success rates.

Figure 16 also illustrates the position and velocity errors for the simulated landing scenario. It can be observed in this figure that the position error remains approximately 10 meters, indicating a minor deviation from the target landing site. Additionally, the velocity error at touchdown reaches 8 cm/s, demonstrating the system's capability to achieve a controlled descent with minimal velocity discrepancies. These error margins highlight the accuracy of the trajectory optimization and control algorithms in guiding the spacecraft to a precise and safe landing.




Figure 15. Lateral view of the trajectory shown in Figure 14. Emphasis is put on the sharp turn the trajectory takes once it switches the target destination from the hovering point to the landing spot. The pyramid has been removed to provide a clearer view of the geometry [23].




Figure 16. Position and velocity evolution plots for the computed landing trajectory. On the left-hand side, the full evolution is shown. On the right-hand side, a detailed view of the last phase is shown.

5. Conclusions and Future Work

A novel methodology has been presented where a fully autonomous GNC system for on-board mission planning is presented. The system relies on visual inputs obtained uniquely from passive sensors (cameras) and the acquired data is fully processed using CNN-based algorithms, which reduce the computational burden of the task, making it feasible to run all of this live on the spacecraft. This methodology, combined with an MPC-based trajectory optimisation block, is meant to provide the means for a spacecraft to learn about a small-body environment, select a landing spot that fits the mission requirements and needs, and design the trajectory leading to this landing sequence; all in an autonomous fashion.

A surveillance or exploration segment is required for missions featuring this methodology where the body of interest is orbited and mapped appropriately. The duration of this phase will depend on the orbit design and the mission objectives. During this phase, images are acquired until the GSM is populated. Illumination conditions play a critical role in how fast the GSM is populated with hazard-related information.

Once the GSM reaches the level of information required, an optimal landing spot is selected based on mission requirements and a distance-from-hazard metric, to reduce as much as possible the chance of landing on an unsafe or hazardous area.

This methodology introduces a new paradigm for missions to small bodies, providing a process pipeline where different components can be plugged in to fit mission specifications and sensor needs. Being able to combine surface observations with pose estimation techniques reinforces greatly the on-board navigation capabilities of any mission and the use of ML-based techniques opens up a great deal of possibilities for autonomous solutions.

However, many questions remain after this initial step towards the objective of a fully autonomous GNC system. For instance, the study of reconnaissance orbits to optimise the population of the GSMs is a topic that could represent a field of research itself. Not only designing the orbit for coverage of the body but to optimise for a hazard evaluation around particular areas of the body would require planning for the ideal trajectory and would couple that with ideal illumination conditions to acquire significant images that could contribute the most to the GSM.

Another aspect would be the effect of uncertainties on the pose estimation block. Studying how sensitive this methodology is to errors in the estimation of the exact position or attitude of the spacecraft would increase its robustness greatly. Depending on the results, a new priority scheme might have to be developed, with the covariance of the observations (or any other metric for the accuracy of the

estimation) as a parameter to rank how reliable a certain observation and its contribution to the GSM is.

In any case, this work paves the way for autonomous missions that are capable of arriving at an unknown body, and exploring it by themselves. Such a paradigm, although still distant, would represent a game-changing step forward for space exploration.

Funding: The authors of this paper would like to acknowledge the funding received from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 813644.

The authors would like to acknowledge too the funding received for APC from Universidad Carlos III de Madrid under the Read & Publish Agreement CRUE-CSIC 2024.



Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Acronyms and abbreviations

The following acronyms and abbreviations are used in this manuscript:

AI	Artificial Intelligence
ART	Autonomous Rendez-vous Transformer
CNN	Convolutional Neural Network
GNC	Guidance, Navigation, and Control
GSM	Global Safety Map
HDA	Hazard Detection and Avoidance
IP	Image Processing
KOZ	Keep-Out Zone
KSPDG	Kerbal Space Program Differential Games
LLM	Large Language Model
LQP	Linear Quadratic optimization Problem
MPC	Model Predictive Control
OSIRIS-REx	Origins, Spectral Interpretation, Resource Identification, and Security–Regolith Explorer
PfS	Pole-from-Silhouette
RAAN	Right Ascension of the Ascending Node
ResNet	Residual Network
SfM	Structure-from-Motion
TaG	Touch-and-Go

References

1. Lauretta, D.S.; Balram-Knutson, S.S.; Beshore, E.; Boynton, W.V.; Drouet d’Aubigny, C.; DellaGiustina, D.N.; Enos, H.L.; Golish, D.R.; Hergenrother, C.W.; Howell, E.S.; Bennett, C.A.; Morton, E.T.; Nolan, M.C.; Rizk, B.; Roper, H.L.; Bartels, A.E.; Bos, B.J.; Dworkin, J.P.; Highsmith, D.E.; Lorenz, D.A.; Lim, L.F.; Mink, R.; Moreau, M.C.; Nuth, J.A.; Reuter, D.C.; Simon, A.A.; Bierhaus, E.B.; Bryan, B.H.; Ballouz, R.; Barnouin, O.S.; Binzel, R.P.; Bottke, W.F.; Hamilton, V.E.; Walsh, K.J.; Chesley, S.R.; Christensen, P.R.; Clark, B.E.; Connolly, H.C.; Crombie, M.K.; Daly, M.G.; Emery, J.P.; McCoy, T.J.; McMahon, J.W.; Scheeres, D.J.; Messenger, S.; Nakamura-Messenger, K.; Richter, K.; Sandford, S.A. OSIRIS-REx: Sample Return from Asteroid (101955) Bennu. *Space Sci Rev* **2017**, *212*, 925–984. doi:10.1007/s11214-017-0405-1.

2. Villa, J.; McMahon, J.; Nesnas, I. *Robust Landmark and Hazard Detection on Small Body Surfaces Using Shadow Imagery*; 2022.
3. Villa, J.; McMahon, J.; Hockman, B.; Nesnas, I. *Autonomous Navigation and Dense Shape Reconstruction Using Stereophotogrammetry at Small Celestial Bodies*; 2022.
4. Arvidson, R.E.; Bell III, J.F.; Bellutta, P.; Cabrol, N.A.; Catalano, J.G.; Cohen, J.; Crumpler, L.S.; Des Marais, D.J.; Estlin, T.A.; Farrand, W.H.; Gellert, R.; Grant, J.A.; Greenberger, R.N.; Guinness, E.A.; Herkenhoff, K.E.; Herman, J.A.; Iagnemma, K.D.; Johnson, J.R.; Klingelhöfer, G.; Li, R.; Lichtenberg, K.A.; Maxwell, S.A.; Ming, D.W.; Morris, R.V.; Rice, M.S.; Ruff, S.W.; Shaw, A.; Siebach, K.L.; de Souza, P.A.; Stroupe, A.W.; Squyres, S.W.; Sullivan, R.J.; Talley, K.P.; Townsend, J.A.; Wang, A.; Wright, J.R.; Yen, A.S. Spirit Mars Rover Mission: Overview and selected results from the northern Home Plate Winter Haven to the side of Scamander crater. *Journal of Geophysical Research: Planets* **2010**, *115*. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2010JE003633>, doi:10.1029/2010JE003633.
5. Welch, R.; Limonadi, D.; Manning, R. Systems engineering the Curiosity Rover: A retrospective. 2013 8th International Conference on System of Systems Engineering, 2013, pp. 70–75. doi:10.1109/SYSoSE.2013.6575245.
6. Arvidson, R.E. Roving on Mars with Opportunity and Curiosity: Terramechanics and Terrain Properties **2015**. pp. 165–173. Publisher: American Society of Civil Engineers, doi:10.1061/9780784479179.019.
7. Glassmeier, K.H.; Boehnhardt, H.; Koschny, D.; Kührt, E.; Richter, I. The Rosetta Mission: Flying Towards the Origin of the Solar System. *Space Sci Rev* **2007**, *128*, 1–21. doi:10.1007/s11214-006-9140-8.
8. Frolov, V.; Faizov, B.; Shakhuro, V.; Sanzharov, V.; Konushin, A.; Galaktionov, V.; Voloboy, A. Image Synthesis Pipeline for CNN-Based Sensing Systems. *Sensors* **2022**, *22*, 2080. Number: 6 Publisher: Multidisciplinary Digital Publishing Institute, doi:10.3390/s22062080.
9. Schubert, C.; Black, K.; Fonseka, D.; Dhir, A.; Deutsch, J.; Dhamani, N.; Martin, G.; Akella, M. A Pipeline for Vision-Based On-Orbit Proximity Operations Using Deep Learning and Synthetic Imagery. 2021 IEEE Aerospace Conference (50100), 2021, pp. 1–15. ISSN: 1095-323X, doi:10.1109/AERO50100.2021.9438232.
10. Pugliatti, M.; Franzese, V.; Topputo, F. Data-Driven Image Processing for Onboard Optical Navigation Around a Binary Asteroid. *Journal of Spacecraft and Rockets* **2022**, *59*, 943–959. Publisher: American Institute of Aeronautics and Astronautics _eprint: <https://doi.org/10.2514/1.A35213>, doi:10.2514/1.A35213.
11. Guffanti, T.; Gammelli, D.; D'Amico, S.; Pavone, M. Transformers for Trajectory Optimization with Application to Spacecraft Rendezvous. 2024 IEEE Aerospace Conference, 2024, pp. 1–13. ISSN: 1095-323X, doi:10.1109/AERO58975.2024.10521334.
12. University Carlos III of Madrid, Leganés, Spain.; Escalante, A.; Sanjurjo Rivo, M.; Ghiglini, P. Bennunet - Applying Machine Learning Techniques for Autonomous Optical Relative Navigation of an Asteroid. Papers of ESA GNC-ICATT 2023. ESA, 2023. doi:10.5270/esa-gnc-icatt-2023-015.
13. Epp, C.D.; Smith, T.B. Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT). 2007 IEEE Aerospace Conference, 2007, pp. 1–7. ISSN: 1095-323X, doi:10.1109/AERO.2007.352724.
14. Parreira, B.; Vasconcelos, J.F.; Montaña, J.; Ramón, J.; Penin, L.F. Hazard Detection and Avoidance in ESA Lunar Lander: Concept and Performance. AIAA Guidance, Navigation, and Control (GNC) Conference; American Institute of Aeronautics and Astronautics: Boston, MA, 2013. doi:10.2514/6.2013-5020.
15. Tomita, K.; Skinner, K.; Iiyama, K.; Jagatia, B.; Nakagawa, T.; Ho, K. Hazard Detection Algorithm for Planetary Landing Using Semantic Segmentation. ASCEND 2020; American Institute of Aeronautics and Astronautics: Virtual Event, 2020. doi:10.2514/6.2020-4150.
16. Peñarroya, P.; Centuori, S.; Sanjurjo, M.; Hermosín, P. A LiDAR-less approach to autonomous hazard detection and avoidance systems based on semantic segmentation. *Celest Mech Dyn Astron* **2023**, *135*, 34. doi:10.1007/s10569-023-10140-9.
17. Baker, D.; McMahon, J. *Limb-Based Shape Modeling: A Demonstration on Itokawa*; 2019.
18. Bandyopadhyay, S.; Villa, J.; Osmundson, A.; Hockman, B.; Morrell, B.; Lubey, D.; Bhaskaran, S.; Bayard, D.; McMahon, J.; Nesnas, I. *Light-Robust Pole-from-Silhouette Algorithm and Visual-Hull Estimation for Autonomous Optical Navigation to an Unknown Small Body*; 2021.
19. Rodriguez-Fernandez, V.; Carrasco, A.; Cheng, J.; Scharf, E.; Siew, P.M.; Linares, R. Language Models are Spacecraft Operators, 2024. arXiv:2404.00413 [physics], doi:10.48550/arXiv.2404.00413.
20. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition, 2015. arXiv:1512.03385 [cs], doi:10.48550/arXiv.1512.03385.

21. Akçay, S.; Kundegorski, M.E.; Devereux, M.; Breckon, T.P. Transfer learning using convolutional neural networks for object classification within X-ray baggage security imagery. 2016 IEEE International Conference on Image Processing (ICIP), 2016, pp. 1057–1061. ISSN: 2381-8549, doi:10.1109/ICIP.2016.7532519.
22. Frekhaug, T.; Sanjurjo-Rivo, M.; Soler, M. Output-Feedback Tube Model Predictive Control for Robust Autonomous Descent and Landing on Asteroids. 2022.
23. Frekhaug, T.A.; Escalante, A.; Sanjurjo-Rivo, M.; Soler, M. Robust Model Predictive Control with Monocular Optical Navigation System for Asteroid Circumnavigation. 21st European Control Conference (ECC), 2023, pp. 1938–1941.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.