(Preprint) AAS 25-173

A SURVEY OF AUTONOMOUS NAVIGATION TECHNIQUES APPLICABLE TO LUNAR SURFACE EXPLORATION

Paul D. McKee*

As humanity returns to the Moon, and more and more attention is being paid to lunar surface operations, there is a greater need than ever for methods of surface navigation. These could be methods of computer-assisted orienteering for astronauts exploring on foot during an Extra-Vehicular Activity (EVA), or methods of solving the Lost-on-the-Moon problem to initialize a crewed or autonomous rover's state estimate. It may also be necessary to process navigation data associated with surface samples or other surface operations *a posteriori* to better understand where that analysis occurred. Autonomous rover operation will also require Hazard Detection and Avoidance (HDA) and terrain-aware pathfinding.

While navigation on the surface of the Moon will likely rely on Earth-based assets such as the Deep Space Network (DSN) or communication with other spacecraft (e.g., LunaNet, LCRNS, pre-deployed moon beacons, a nearby lander) it may be necessary to navigate in a loss-of-communication scenario. This paper analyzes the methods of surface navigation used on other celestial bodies, such as those used during the Apollo missions and autonomous exploration of Mars, as well as novel methods which have been studied but not yet implemented which may prove useful. It is shown that the navigator has myriad options when processing data from an Inertial Measurement Unit (IMU), a star tracker, (rover) wheel encoders, optical cameras, and LIght Detection and Ranging (LIDAR) sensors. The intention of this paper is to provide a broad overview of what has been done and what could be done, to aid those designing vehicles and/or missions to the lunar surface.

INTRODUCTION

Getting lost can ruin the whole trip. This is doubly true when traversing the surface of another world. As nations send spacecraft to the Moon at an ever more rapid pace, and as the Artemis program promises to establish a sustained human presence on the Moon, the need for reliable lunar surface navigation has never been greater. This analysis will focus on localization, or position estimation, which will be necessary for the safe completion of any mission, crewed or uncrewed. Navigation on the lunar surface will almost certainly leverage assets on the surface of the Earth, in Earth orbit, in lunar orbit, or other assets on the surface of the Moon. These could include the Deep Space Network (DSN),[1] the Global Positioning System (GPS),[2,3] spacecraft in lunar orbit following the LunaNet protocol,[4] or pre-deployed "moon beacons".[5,6] However, autonomous navigation capabilities on the lunar surface would be helpful in the case of loss-of-communication (a genuine concern at the south pole of the Moon), to relieve demand on other communications assets, or to "tag" locations of surface operations or sample acquisition using concurrent navigation data. This analysis will consider the problem of lunar surface localization for a vehicle (crewed or uncrewed) or a crew-member on an Extra-Vehicular Activity (EVA) using only the navigation instruments that can reasonably be brought along.

*GN&C Engineer, Aeroscience and Flight Mechanics Division, NASA-JSC, 2101 NASA Pkwy, Houston, TX 77058

Within the broader category of navigation there are several smaller problems that can be considered. The first is global navigation, estimating position on the surface of the Moon in some global frame (i.e., finding latitude and longitude). The second is local navigation, finding how far one has traveled and in what direction with respect to a starting point or some other asset. The third is Hazard Detection and Avoidance (HDA), recognizing potential dangers and planning a safe traverse path that avoids these. Several navigation sensors could be employed to achieve these goals, including an IMU (consisting of accelerometers and gyroscopes), a star tracker, wheel encoders (in the case of a wheeled vehicle), optical navigation cameras, or a LIght Detection and Ranging (LIDAR) instrument. This analysis will consider several navigation methods and algorithms, and discuss which of these instruments would be necessary for each.

There is much to be learned from past and current missions to other celestial bodies, and the navigation methods utilized to date. This paper includes three such case studies: the Apollo Lunar Roving Vehicle (LRV, a.k.a. "Moon Buggy"), the Mars rovers designed and operated by NASA's Jet Propulsion Laboratory (JPL), and the Volatiles Investigating Polar Exploration Rover (VIPER) that was designed by NASA Ames Research Center and is under construction at NASA Johnson Space Center (JSC). The navigation hardware, algorithms, and Concept of Operations (ConOps) for each vehicle will be explored. The rest of the paper consists of a brief primer on several methods of local and global surface navigation and HDA which have been implemented or studied, as well as comments on the applicability of each to future missions and vehicles.

CASE STUDY: APOLLO LUNAR ROVING VEHICLE

Engineers working on the LRV for the later Apollo missions faced a unique challenge. They had to write the requirements for the first-ever lunar surface navigation system, then actually build and implement said system, subject to immense cost and schedule pressure. Proposals for methods of lunar surface navigation were published as far back as the mid-1960's and some introduced techniques that have not flown even to this day. The navigation system for the LRV was designed by Boeing and Bellcomm Inc. with assistance from NASA Marshall Space Flight Center (MSFC) and from the NASA Manned Spacecraft Center (MSC)—now the Lyndon B. Johnson Space Center. Bill Tindall, in one of his famed "Tindallgrams," concluded that an intricate navigation system was not needed, nor was there time to develop one.[7] Whatever flew must be as simple as possible. There would be no global navigation on the LRV, as it would not be necessary for such a short mission. Instead, "the crew's eyeballs and map will have the highest priority." The LRV would provide "an indication of their heading and distance travelled." Finally, the LRV provided no HDA capabilities—the crew would take on that responsibility as well. In a technical memo about the LRV navigation system published in 1970, it was clearly stated that "the navigation system is only a convenience item, not mandatory, and that mission rules should not require the system to be operable either to begin or continue a sortie".[8]

One of the most popular and promising recommendations for LRV global navigation (which did not fly) was a "stellar-inertial system" which would utilize a star tracker and a method of measuring the local gravity vector, such as an IMU or multi-axis pendulum system.[9,10] One of the limiting factors would be the understanding the Moon's gravitational field, as errors in the local "down" vector cascade into global position errors. It was also suggested that a rover could be tracked using Earth-based measurements[9] and this was actually done with LRV after the fact.[11] Both the LRV and the LM had S-band transmitters which were tracked by the Goddard Space Flight Tracking Network. Using Very-Long Baseline Interferometry (VLBI), the rover's position was found *a posteriori* to

within 25 meters meters of error.[11] Another global navigation suggestion was that the crew navigate visually using landmarks or pre-placed radio-frequency (RF) beacons, which present a challenge of their own.[9] It was noted then (and is still relevant now) that a lunar rover cannot rely on any sort of compass, either magnetic or gyroscopic in nature, as the Moon's magnetic field is too weak and its rate of spin is too slow for either to be useful navigation observables.[10] Rather, heading information must come from an IMU, which is subject to drift and must be realigned periodically using external attitude measurements, such as stellar measurements or Sun angles.

Throughout the Apollo program, surface global navigation was to be performed by the crew using their eyes and a map. Simple orienteering approaches were used, similar to those used by hikers on Earth. These methods famously fell short on the Apollo 14 mission, when the crew became disoriented when looking for Cone crater.[12] They consistently overestimated the distance they had traveled, and became exhausted climbing crater rim slopes. It was noted that the lunar surface "plays tricks on" the human eye, due to terrain self-similarity, extreme lighting conditions, and no true sense of scale (is that boulder close and near, or huge and far away?). A joint crew-ground decision was made to abandon their search for Cone crater, due to risk of overexertion. Later analysis showed that the crew were within 30 meters of their destination when this decision was made. Thankfully, there were no major shortcomings of orienteering on the rover missions.

Hazard detection and avoidance was also left up to the crew. The primary hazards on the lunar surface are rocks and craters. Astronauts reported that rocks were easy enough to sight at a distance and maneuver around, but craters could sneak up on an otherwise vigilant driver. More than once, evasive action was required to avoid driving into a crater.[12] The sloped wall of a crater could cause the vehicle to roll over. Worse, the crew and vehicle could become stuck in a steep-walled crater from which they couldn't climb out. It would be prudent to equip a modern lunar rover with sensors and algorithms to detect rocks and craters and alert the crew to potential hazards in real-time.

The local navigation system that was implemented on the LRV was concisely described by Bill Tindall as "a compass and an odometer." In reality, the heading of the rover was tracked by an IMU which had to be aligned periodically using an external reference. Knowing the time history of the speed and the heading of the vehicle, it is straightforward to integrate and determine how far it has travelled and in what direction. The maximum planned distance from the Lunar Module was 9.5km, and it was believed that the LM should be plainly visible out to 1km and probably out to 3km.[8] The goal, then, was to provide the crew with enough information to drive back to a point within visual range of the LM. The center console of the LRV displayed the distance traveled (which was inversely proportional to the remaining battery charge), the bearing and range to the LM, and the speed, as illustrated in Fig. 1. The crew commented that these numbers were remarkably accurate, and that it was easy to drive back to the LM from the end-point of their EVA.[12]

Each of the four rover wheels had an odometer, and concurrent signal between at least three was required to record one revolution. This logic was used to combat errors due to one of the wheels slipping. Preliminary analysis indicated that range error due to wheel slip would likely contribute much less to overall position error than bearing error due to gyroscope misalignment.[8] Gyro realignment was performed at the end of each station of an EVA, just before the crew drove to the next destination. A Sun shadow device (seen in Fig. 1) and a vehicle pitch and roll indicator (not shown) were used in conjunction with ephemeris data on the location of the Earth, Moon, and Sun. The Sun shadow angle, as measured by the device on the LRV console, is a function of the Sun azimuth with respect to lunar north, the Sun elevation above the horizon, the LRV heading with respect to lunar north, the relative heading of the LRV with respect to the Sun, the LRV pitch and roll

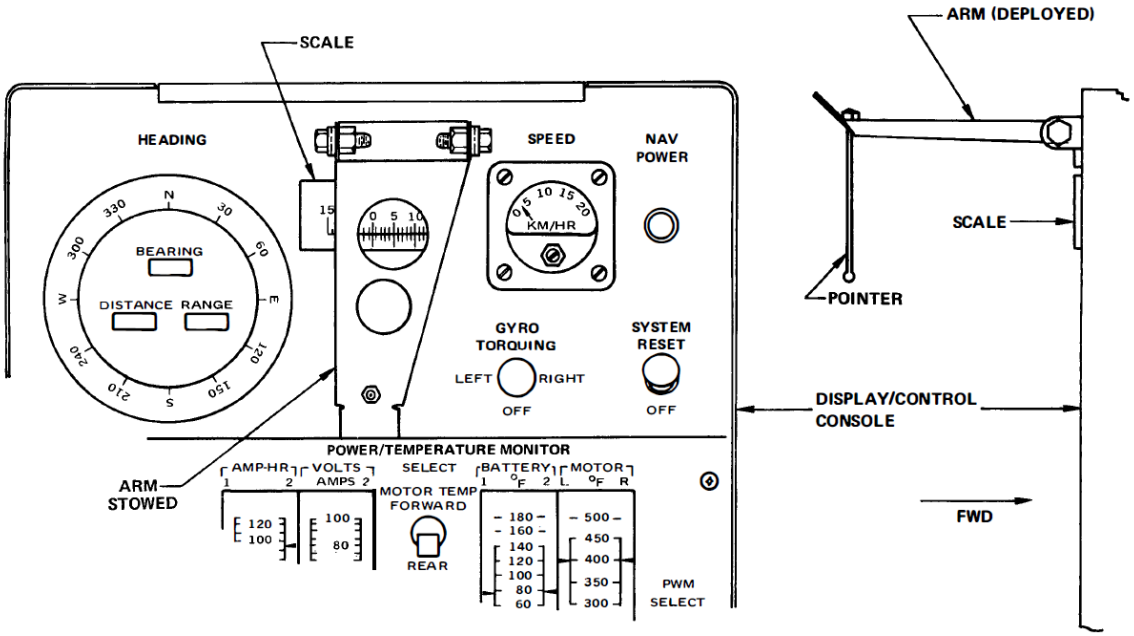

Figure 1. The Apollo LRV center console displayed the range and bearing from the rover's current position back to the LM. Reprinted from NASA TM-70-2014-8[8]

angles, and the Sun shadow device size and geometry. Some of this information could be read by the crew and the rest was book-kept by mission control. A calculation was performed on the ground to find the LRV heading, and the crew could realign the IMU before proceeding with their EVA. This method of IMU realignment sacrificed autonomy for onboard simplicity. Radio communication with Earth was required, as were ground-based computers, but this could be accomplished with a very simple analog device. Something similar could be performed today with a (likely zenith-facing) star tracker in a completely autonomous manner.

The navigation system of the LRV was studied and scrutinized after-the-fact and was determined to have performed very well. Bellcomm Inc. computed IMU alignment headings and it was determined that the gyro virtually did not drift on Apollo 15.[13] The vehicle accuracy requirements were as follows: the bearing to the LM had to be known to $\pm 6\,\mathrm{deg}$, the range to the LM had to be known to ± 600 meters, and the total distance traveled had to be known to $\pm 2\%$. A comprehensive LRV navigation system performance review determined that the system performed within specifications on Apollo 15, 16, and 17.[14]

The LRV navigation system performed admirably and was truly a triumph of simplicity. However, a modern lunar surface rover could accomplish much more. Improvements in navigation and autonomy capabilities over the last decades could yield tremendous gains in lunar surface navigation performance. This could be accomplished with relatively inexpensive navigation equipment such as optical cameras. There are several lessons to be gleaned from studying Apollo surface navigation. First, the lunar surface can play tricks on the human eye. Second, the crew should be furnished with enough information to drive to a point within visual range of lander at all times. Third, the crew need not be supplied any excess information. The navigation system should exist to facilitate exploration, and should never pose an undue burden on the crew's time or attention.

CASE STUDY: MARS ROVERS, PAST AND PRESENT

Traversing the surface of Mars requires a much greater degree of autonomy than traversing the surface of the Moon. Rovers on Mars have an extremely high latency with operators on Earth, compared to the much lower latency between the rover and its astronaut operators. As a result, real-time operation is impossible, shifting the priority towards extreme caution. "Slow and steady wins the race" is the rule when it comes to robotic rover operation. In general, the ConOps is to roll forward a short distance, pause, take several measurements for navigation and hazard detection, then send this information back to Earth for one or more human operators to consider before proceeding. Over time, more and more decision-making power has been placed in the hands of on-board autonomous path planning algorithms, but the navigation procedure is still very much human-in-the-loop.

A great deal of pioneering work has been done on image processing, computer vision, hazard detection, path planning, and autonomy in support of Mars surface exploration. Designers of a future rover on any celestial body would do well to study the hardware and algorithms implemented on Mars rovers. At the time of this writing there have been five rovers sent to the Martian surface over the span of four missions, and two of those are still in operation. Fig. 2 (left) shows the mock-ups of the first three Mars rovers—Sojourner (which is roughly the size of a remote-controlled toy car), Spirit and/or Opportunity (which are roughly the size of a go-kart), and Curiosity (which is roughly the size of a Sport Utility Vehicle). Fig. 2 (right) is a photo taken by the Ingenuity Helicopter on the Mars 2020 mission showing its own shadow and the tracks left by the Perseverance rover.

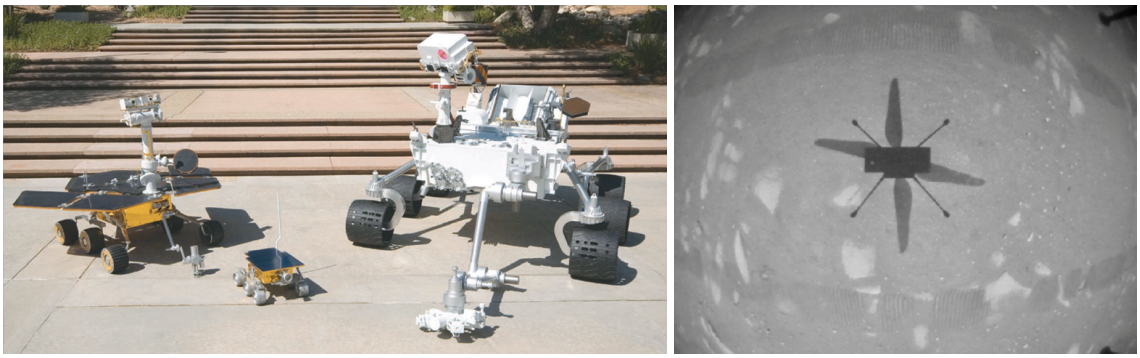

Figure 2. Left: Family of Mars Rover mockups. Sojourner is in the foreground, Mars Exploration Rover (Spirit and Opportunity) is on the left, and Mars Science Laboratory rover (Curiosity) is on the right (image credit: NASA/JPL). Right: Nadir-facing image taken by the Ingenuity helicopter on the Mars 2020 mission shows its own shadow and the tracks of the Perseverance rover (image credit: NASA/JPL).

The Mars Pathfinder mission landed on the Martian surface on July 4th, 1997 and deployed the Sojourner rover shortly thereafter.[15] This was the first ever wheeled vehicle on Mars, and was relatively simple and unassuming compared to its far more complex successors. The rover was originally only meant to function for 7 to 30 sols (1 sol = 1 Martian day $\approx$ 24hr 39min), but operated for 83 sols before the end of the mission.[15,16] The rover explored an area tens of meters in radius around the Pathfinder lander, and would work in conjunction with the Pathfinder lander for navigation—the two often took images of one another.

Sojourner was equipped with stereo monochrome navigation cameras and five infrared laser stripers.[17,18] It also had a single aft-facing color camera which was not used for navigation.[18]

Sojourner had wheel encoders, gyros, and accelerometers[15,19] as well as contact bump sensors.[17] Global navigation was not necessary for the Sojourner rover—it never strayed beyond visual range of the Pathfinder lander. Local navigation was similar in principle to what was done on the Apollo LRV. Six wheel encoders were averaged and then integrated to determine distance traveled, and attitude rate data from a gyro was integrated to determine direction traveled.[15,19] The gyro-based heading reference system was subject to a drift of about 13 degrees per sol, meaning the rover did not always reach its desired waypoint within acceptable bounds. The gyros were realigned after every traverse using the local gravity vector as measured by the accelerometers.[19] The true position of the rover was found after each traverse by a team of human operators studying images taken both by the rover and of the rover (by Pathfinder).[18]

The hazard detection and avoidance system of the Sojourner rover was purely reactive in nature. There was no on-board terrain map and no memory of past hazards encountered.[17] The rover could detect dangerous terrain directly in front of it and decide to turn in place and try a different direction. Laser stripers projected a predictable pattern onto the Martian surface in front of the rover, and stereo cameras captured and processed images of the striped terrain. Deviations in the images from the known pattern indicated the presence of rocks or other terrain hazards. Between this vision-based HDA system, the accelerometers and gyroscopes (which could detect excessive rover tilt), and contact bump sensors, the rover could effectively avoid rocks, drop-offs, hazardous slopes, and collisions with undetected rocks.[15,19] The rover would drive towards a specified waypoint, stop every seven centimeters to perform HDA, and either proceed or turn in place. It repeated this pattern until it (thought it) reached its destination or some maximum amount of time had passed. Even with comparatively barebones navigation and HDA capabilities, Sojourner performed above and beyond expectations.

The Mars Exploration Rover (MER) mission saw twin rovers delivered to the Martian surface on nearly opposite sides of the planet. Spirit began operation on January 4th, 2004, and Opportunity began operation on January 25th, 2004. The rovers were originally designed to operate for 90 sols each. However, Spirit operated for 2208 sols and Opportunity for 5352 sols. These rovers traveled many kilometers from their landing sites, necessitating a means of global navigation. Spirit and Opportunity were essentially twins with the same hardware and software. As such, they will henceforth be referred to collectively as "MER" for the sake of brevity.

Each MER rover had six engineering cameras: a stereo pair of navigation cameras (NavCams) on a point-able mast and a forward-facing and aft-facing pair of body-mounted hazard cameras (HazCams) for hazard detection.[20] There was also a panoramic camera (PanCam) which was primarily intended for science, but whose panoramas proved useful for global navigation. The NavCams could also pivot to take a 360-degree panorama.[20,21] The MER rover had wheel encoders, suspension potentiometers, and an IMU to keep track of vehicle motion.[21] Local navigation was carried out with an IMU and wheel odometry,[22] similar to what was done with Sojourner, but the method of realigning the gyros was much more like the LRVs. The NavCams took take images of the Sun and combining this information with ephemeris data to compute vehicle attitude.[17,21] Global navigation was accomplished using panoramic images which could be taken using either the PanCam or the NavCams. At the beginning of the mission, a "site frame" was defined and the rover navigated relative to that site frame. After a long enough traverse, a new site frame was periodically defined. Whenever this happened, the rover stopped and took a new panorama and sent it back to Earth. Human operators could determine the pose of the rover with respect to known Martian landmarks, thus providing a global pose estimate.[20]

Hazard detection and avoidance took a giant leap from Sojourner to MER. Stereo images taken by the NavCams were used to produce a 3D point cloud of the terrain in front of the rover.[17] That point cloud could then be analyzed for hazards, finding patches onto which a plane the size of the total rover footprint would or would not fit nicely. A "goodness map" of local terrain was then produced, which allowed for path planning.[17,21,22] The single largest improvement in local navigation on the MER rover, which became standard on all future Mars rovers, was visual odometry.[17,21–24] Visual odometry is accomplished by comparing "before" and "after" images of the same terrain and tracking the apparent motion of image feature points. Visual odometry is a computationally expensive image processing technique, and thus was originally planned to be used only when absolutely necessary.[17] On sol 446, Opportunity was commanded to drive 50 meters but failed to clear "Purgatory Ripple"—driving only two meters and literally spinning its wheels. This issue was identified using visual odometry, which was also instrumental in navigating around this obstacle.[17,23] Both visual odometry and stereo-image-based surface hazard detection were pioneered on the MER mission and became standard practice on future Mars rovers.

The Mars Science Laboratory (MSL) mission landed on August 6th, 2012, and deployed the Curiosity rover,[25] which is still operational at the time of this writing. Curiosity and has traveled over 30km and operated for over 4400 sols so far. This vastly exceeds its original planned operational lifetime of one Martian year (668.6 sols, just under two Earth years). Like the Spirit and Opportunity rovers, Curiosity would be nomadic rather than staying near its landing site. Curiosity has four NavCams and eight HazCams—two pairs of forward-facing HazCams and two pairs of aft-facing HazCams—which were all essentially duplicates of the cameras used on Spirit and Opportunity. The NavCams formed two redundant stereo pairs on the rover mast, which could pan and tilt to better orient the cameras. The forward-facing HazCams could each see both front wheels within a single image, but the aft-facing HazCams could each only see one wheel due to the configuration of the cameras, wheels, and the rear-mounted Radioisotope Thermoelectric Generator (RTG) which served as the rover's power source.[26] The rover also had an IMU and wheel encoders to help keep track of relative motion.[25]

Local navigation on Curiosity was done in much the same way as it was done on the MER rovers. The IMU and wheel encoders were integrated to find the change in rover pose over a period of time.[25,27] Sun bearing data from the NavCams could be combined with accelerometer measurements and known ephemeris data for an attitude update.[28] Visual odometry was considered to be a main feature of the navigation arsenal. Images perpendicular to the direction of vehicle motion were taken to maximize pose observability.[27] Panoramic images were taken by the MastCam (a science instrument) and/or NavCams at the end of each traverse[27] and could be used for global navigation with the assistance of human operators on Earth. Hazard detection and avoidance for the Curiosity rover also depended on stereo image processing to create point clouds which were analyzed for hazards and feasible traverse paths. More of the processing and path planning was done on-board Curiosity than was done on the Spirit or Opportunity rovers, marking a step in the direction of greater rover autonomy.[27] The MSL mission showcased a similar navigation ConOps to the MER mission, but with several incremental improvements.

The Mars 2020 mission, which launched on July 30th, 2020, and landed on February 18th, 2021, saw the deployment of the Perseverance rover and the Ingenuity helicopter—the two newest members of the Mars rover family.[29] The mission was originally planned to last approximately one Martian year. The Perseverance rover has exceeded that lifetime and has been operating on Mars for over 1100 sols at the time of this writing. The navigation ConOps of the Perseverance rover closely

mirrors those of MER and MSL: to do point-to-point exploration rather than lingering around the landing site for the entire mission. Perseverance, nicknamed "Percy", has been tasked with collecting and containing surface samples and leaving these at strategic locations as the first phase of a Mars Sample Return mission.

Perseverance has one stereo pair of mast-mounted NavCams, two pairs of forward-facing HazCams, and one pair of aft-facing HazCams. These cameras represent a significant upgrade over those of MER and MSL.[30] Percy is equipped with a Field-Programmable Gate Array (FPGA) for processing stereo images.[31] As a result, image processing and computer vision feature much more heavily in the Perseverance ConOps. Percy can (1) acquire and process stereo images, (2) analyze terrain and select a best path forward, and (3) execute drive commands all in parallel.[32] Perseverance can do this all while driving at nearly its top speed of 4.4 cm/s.[31] Owing to a larger field-of-view (FOV), Perseverance can image an area wider than its own footprint without slewing its cameras.[31] The path planning algorithm is also more complex. Rather than computing if a rover-sized plane would fit on the terrain, it can use a rover shape model to test and approve regions of terrain which will accommodate the wheels and not scrape the rover underside, allowing for more adventurous roving.[31] Perseverance is sometimes trusted to navigate terrain that no human has ever seen—operators on Earth will execute commands prompting the rover to drive to a location that is outside of the FOV of any image taken so far, and trust that the rover will plan its own safe path to get there.[31]

The Mars 2020 mission also includes the Ingenuity helicopter,[33] which demonstrated the first powered flight on another planet. At the time of this writing, the Ingenuity helicopter has recently stopped flying. A traverse over a "visually bland" region of terrain proved treacherous for the on-board visual odometry. The helicopter suffered a hard-landing in which multiple rotor blades were damaged. Fortunately, the computer, cameras, and communication systems are still operational. Ingenuity flew 72 times, significantly outperforming its original planned operational lifespan: five flights of up to 90 seconds each. Ingenuity was equipped with a primary and backup IMU, an inclinometer, a nadir-pointed LIDAR altimeter, and a nadir-pointed monochrome camera, as well as an FPGA for data processing.[34] Due to hardware and software constraints, only the bare minimum amount of navigation was done onboard Ingenuity. Global navigation was done by humans on Earth, analyzing pre-landing imagery from the previous flight to determine the pose of Ingenuity before the start of the next flight.[35] Hazard avoidance was also done by humans, who studied orbital images, selected a safe landing site, and defined its position pre-flight in terms of the helicopter's current pose.[35]

Most of the on-board processing was done towards local navigation. The IMU reading was integrated over the course of each flight to estimate the distance and direction traveled in three dimensions. This IMU was aligned before each flight using readings from the inclinometer. Some gyro drift was accepted due to the short nature of the flights—only 90 seconds in duration.[34] Visual odometry was performed at 30Hz using the mono camera and the FPGA. Once none of the original features from the first image can be found in the latest image, a new base frame was defined and the process began anew.[35] Visual odometry with a mono image taken in flight more closely resembles that done by orbiters or landers, as described by Christian.[36] Performing visual odometry with a single image introduces a range ambiguity necessitating some other instrument like an altimeter. However, little can be done if an insufficient number of image features are detected for several images in a row. The story of Ingenuity serves to underscore the limitations of visual odometry and the necessity of redundant navigation methods.

The Mars rovers have demonstrated that image-based navigation and hazard detection is practical on another celestial body. To date, global navigation on the surface of Mars has consisted of taking panoramic images and sending them to Earth for analysis by human operators with access to orbital images. Local navigation on Mars has consisted of integrating IMU and wheel odometer data, updating attitude estimates using images of the Sun, and stereo visual odometry. Rover HDA has consisted of stereo image processing, 3D point cloud reconstruction, and path planning using knowledge of the rover geometry. One would expect the next generation of rovers to be equipped with a bevy of cameras for navigation and hazard detection, as well as an IMU, wheel odometers, and a dedicated computer for image processing.

CASE STUDY: VOLATILES INVESTIGATING POLAR EXPLORATION ROVER (VIPER)

The Volatiles Investigating Polar Exploration Rover (VIPER)[37–39] will be NASA's first uncrewed lunar rover. At the time of this writing, the future of VIPER remains unsure after the plan to cancel the program was announced in July 2024. However, work at Ames and JSC continues undeterred in the hope that some new path to the Moon can be negotiated. The purpose of the rover is to investigate the nature and distribution volatile resources on or just below the lunar surface. The presence of water ice has been detected near the lunar south pole, in and around the Permanently Shadowed Regions (PSRs). Future crewed Artemis missions will also target this region for human exploration. VIPER will hopefully serve as a pathfinder mission to study the presence of water, which might be considered for possible In-Situ Resource Utilization (ISRU).

VIPER is a low-cost and high-risk mission when compared to the Mars rovers. Fortunately, VIPER is close enough to Earth for near-real-time operation with humans in the loop. The light delay from Earth to the Moon is just over a second, compared to several minutes to Mars. However, VIPER requires Direct-To-Earth (DTE) communication and regular recharge opportunities throughout the operational life of its solar panels. As such, it must avoid Solar shadows as well as "radio shadows" where lunar terrain blocks signals to and from Earth. Thus, VIPER must adopt a different ConOps than the Mars rovers, and this affects the navigation and path planning strategies for the mission. VIPER is equipped with multiple scientific instruments to drill into the surface and analyze the chemical composition of the regolith. In addition, VIPER is equipped with multiple instruments for navigation,[40] as seen in Fig. 3. VIPER has wheel and suspension resolvers (providing the same navigation data as encoders), as well as an IMU. It has a zenith-facing star tracker, placed in that orientation to maximize the angular distance away from the Sun (which will be near the horizon). On its steerable mast is a stereo pair of navigation cameras (NavCams). The aft of the rover has a stereo pair of hazard cameras (AftCams) which aid in driving and path-planning in the aft direction. There is also a hazard camera (HazCam) in each wheel well.[39] This arrangement of instruments and cameras closely mirrors what has been used in the past on the Mars rovers.

The mast on the VIPER rover can pivot to take a full 360-degree panorama of images. Each stereo pair of images can be used to produce a point cloud which can be compared to a Digital Elevation Map (DEM) of the landing and operations site. Alternatively, surface panoramas can be compared with orbital images of the surface to estimate the rover's pose. This analysis is done on Earth with humans in the loop. VIPER requires continuous X-band communication with Earth via the Deep Space Network (DSN), and enters contingency mode if this communication is ever lost.[41] The tried-and-true method of integrating wheel odometer and IMU data to estimate the vehicle's change in position and heading over time will again be used on VIPER. The IMU drift can be accounted for by taking external measurements with its star tracker.[39] This is in contrast to the Mars rovers, which

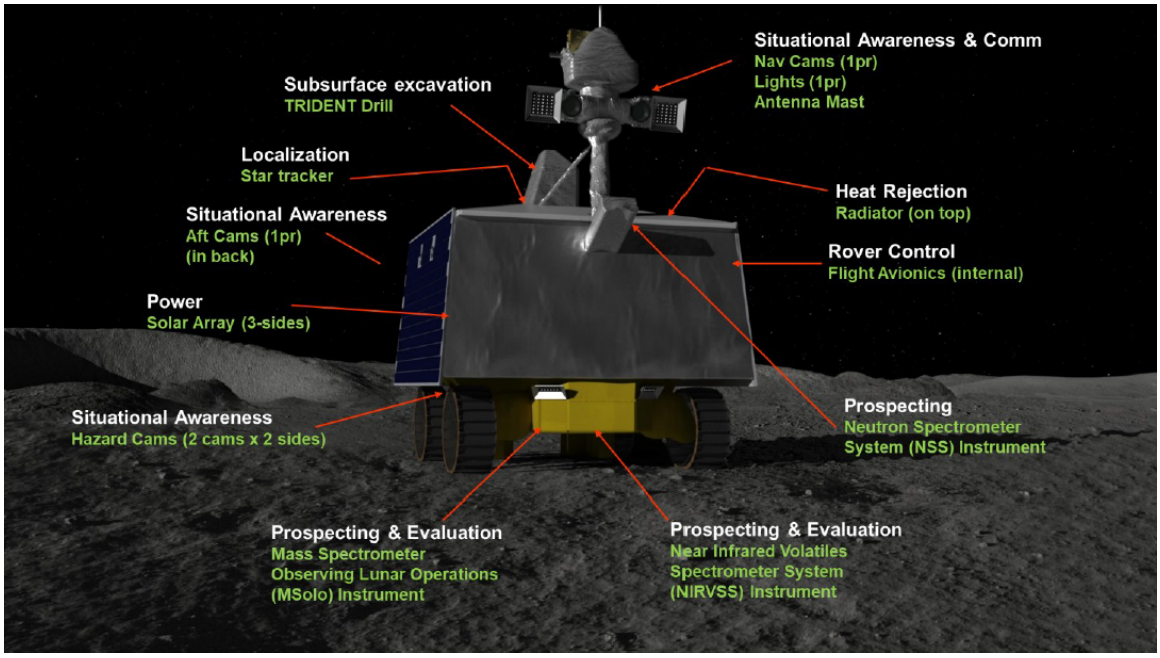

Figure 3. Concept art of the VIPER rover exploring the lunar south pole region. VIPER is equipped with several cameras and a star tracker (among other instruments) for navigation. Reprinted from VIPER Proposal Information Paper (PIP) document[39] with permission from the author.

pointed their NavCams at the Sun for a solar line-of-sight (LOS) measurement, and to the Apollo LRV which used a Sun shadow device (essentially a solar compass). Additionally, images from the navigation cameras can be used for visual odometry.

VIPER must avoid rocks, slopes, and craters, and must also be careful solar shadows which will affect power generation, and radio shadows which interrupt communication with Earth.[42] Because VIPER will be exploring near the lunar south pole, the Sun and the Earth will be along the horizon for the duration of the mission, and these shadows will move at speeds of up to 1cm/s, which is on the same order of magnitude as VIPER's traverse speed. Thus, HDA and path planning to avoid shadows will be a dynamic process with humans in the loop, which would be impossible on Mars due to high latency. The situation on the Moon presents a multi-objective optimization problem, as some of the sites of highest scientific interest are closer to the PSRs, and thus present greater hazards.

After studying the navigation hardware on the Mars rovers and the VIPER lunar rover, it seems that the navigation hardware for rovers has converged. Global navigation is accomplished by using mast-mounted stereo cameras that take panoramic images which are sent back to Earth for processing. Navigation relative to the last known point can be accomplished either through integrating wheel odometer and IMU data—and accounting for IMU drift with some external attitude measurement of the Sun or stars—or through visual odometry using the NavCams. Depending on the nature of the mission, vision-based HDA can be accomplished autonomously or with humans in the loop. Future rovers sent to the surface of the Moon might be expected to function autonomously between crewed missions. The VIPER mission could serve as a template for the navigation strategies and ConOps of such a rover.

INERTIAL NAVIGATION AND WHEEL ODOMETRY

Every rover that has driven on another celestial body has used inertial navigation and wheel odometry. These tools are extremely reliable, well-tested, well-understood, and have a rich flight heritage. The Apollo LRV had a barebones navigation system, but it included gyros and wheel odometers.[8] This method was used on every Mars rover to date[17] and is planned to fly on the VIPER mission.[39] This has also been studied in the context of terrestrial rover navigation and autonomous vehicles.[43–46] This technique is sometimes called "dead reckoning" as it involves propagating the vehicle's state using only internal measurements. However, it can and should be augmented using external measurements for best performance.

The essence of inertial navigation and wheel odometry is integrating velocity to find position. Measurements of inertial attitude from an IMU provide direction, and measurements from wheel encoders provide distance travelled, as illustrated in Fig. 4. There are two major issues: incorrect direction estimates due to gyro drift and incorrect distance estimates due to wheel slip. The solution to gyro drift is to periodically take an external measurement to find the vehicle's true attitude. The Apollo LRV used a sundial device to estimate the Sun direction and ultimately deduce the vehicle's heading with respect to lunar north.[8] The Mars rovers take NavCam images of the Sun and use the apparent Sun LOS direction to update their IMUs.[17] The VIPER rover will accomplish this using a zenith-facing star tracker for attitude estimation.[39] One must carefully plan the frequency of these external measurements around the mission ConOps (time allocated for navigation) and sensor quality—particularly gyro drift rate.

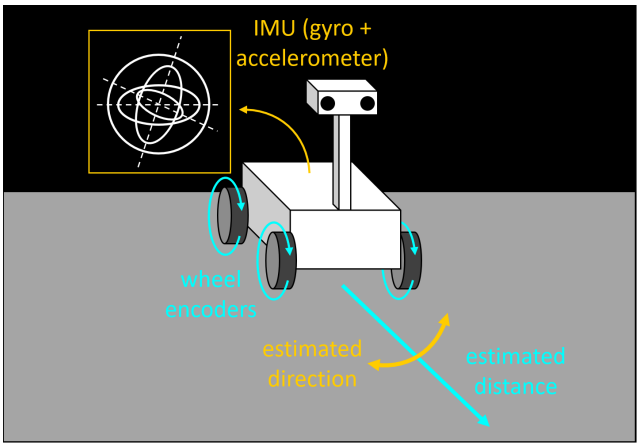

Figure 4. Attitude rate data from an IMU and wheel encoder data can be integrated to estimate distance and direction traveled from some starting point.

The other major issue to contend with is wheel slip. Wheel odometers count the number of revolutions of a wheel, not how much ground has passed underneath it. The Apollo LRV required a concurrent signal from three of its four encoders before it registered one revolution. The Sojourner rover relied on images taken by the nearby Pathfinder lander to verify its dead reckoning solution. Current rovers do (and future rovers almost certainly will) rely on visual odometry as a backup to "old-fashioned" odometry. Alternatively, an accurate model of the nonlinear dynamics of wheel-surface interaction allows for wheel slip prediction and compensation.[43–46]

VISUAL ODOMETRY

Visual odometry estimates the distance a vehicle or observer has traveled using information in digital images. Measuring the apparent motion of image feature points—which correspond to static real-world points—enables the estimation of the change in observer pose. This can be done using a stereo camera pair or a mono camera coupled with some some range measuring device such as a laser range finder or radar altimeter. visual odometry was ran on the Mars Exploration Rovers as an "extra credit" algorithm that ultimately helped to get Opportunity out of a tricky situation, and became critical to the mission.[17,21–24] Since then, visual odometry has become standard on Mars rovers.[25,27,31]

Visual odometry relies on identifying and correlating salient feature points between images. Scale Invariant Feature Transform (SIFT) features[47,48] are well-known, though there are many others (e.g., SURF, BRISK, ORB, KAZE). The "best" feature type is situation-dependent.[49,50] Feature points common to stereo image pairs, or features in a mono image with some external range information, can be localized in 3D space. If these same features are then identified in later images, the change in rover pose between images can be estimated. This is illustrated at a high-level in Fig. 5. Images generally have an abundance of features, so the risk of false match is high. Typically an outlier rejection algorithm such as Random Sample Consensus (RANSAC)[51] is used to keep only feature points that indicate the same change in observer pose.

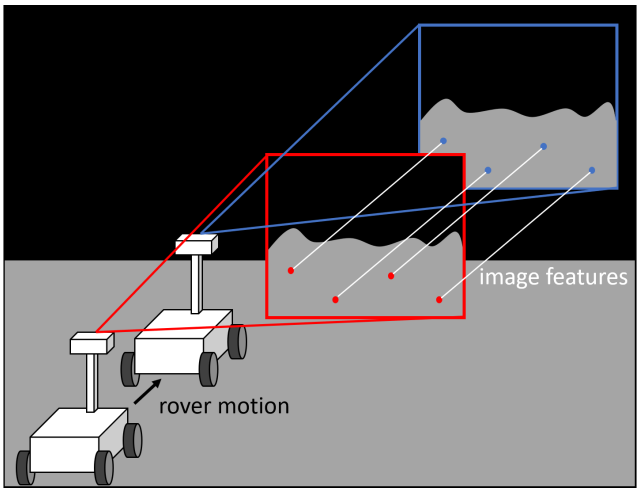

Figure 5. The apparent motion of static terrain features between successive images gives information about the true motion of the observer.

At the time of this writing, visual odometry is operating on the Curiosity rover[27] and Perseverance rover.[31] Maimone et. al.[23] go into some detail on the visual odometry algorithm that flew on the MER rovers and became an increasingly crucial part of the mission. The image processing in visual odometry can be done on board the rover, or it can be done on the ground in post processing.[52] However, if the image processing is meant to be done in anything approximating real-time on the rover, there will need to be a dedicated FPGA or GPU.[31,53] The Perseverance rover has an FPGA for navigation image processing, allowing it to navigate in real-time while driving at its full speed—over 4cm/s.

SIMULTANEOUS LOCALIZATION AND MAPPING

The problem of Simultaneous Localization and Mapping (SLAM) consists of generating a 3D map of the local environment from information available in digital images, then estimating the pose of the observer relative to this map.[54,55] This requires some manner of odometry between image acquisition, as well as the ability to select feature points which can be readily identified and correlated between images. This method is attractive because it requires no *a priori* information about the local environment or observer pose, and this topic is currently an area of rich and active research.

Traditional SLAM can leverage information from cameras, as well as laser range sensors, IMUs, GPS receivers, and wheel encoders. A newer method of SLAM requires only visual information and is referred to as Visual-SLAM (VSLAM).[56,57] One of the crucial considerations for both SLAM and VSLAM is loop closure—the observer must return to (approximately) their original pose and look at some of the initial feature points once again to complete a "loop." This crystallizes the 3D model generated thus far and accounts for odometry errors (e.g., wheel slip, gyro drift). Without loop closure, VSLAM is simply visual odometry. The VSLAM problem is complementary to another similar problem: Structure From Motion (SFM).[58,59] The SFM problem also seeks to generate a 3D model of the environment from information in a series of successive images, but more emphasis is placed on the accuracy of the model and less emphasis is placed on the observer navigation problem.

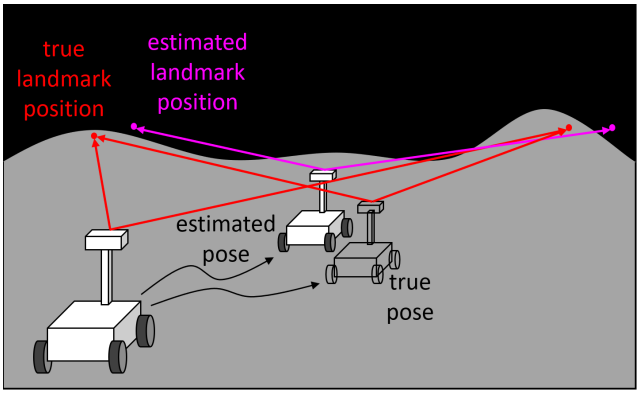

Figure 6. Observer pose estimation and landmark localization can be done concurrently, generating and refining a map over time. Error will accumulate and must be accounted for by returning to the observer's starting point—called "loop closure."

SLAM and VSLAM appear attractive for use on rovers because they require no specialized sensor hardware, and there is a wealth of research and development that had already been done on these topics for terrestrial applications. SLAM requires that a traverse double back on itself at some point, and SLAM-based state data will be of dubious quality until loop closure has been achieved. Fig. 6 illustrates the tendency of SLAM to produce simultaneous incorrect estimates of observer pose and landmark locations which is (often) cleaned up at loop closure. A great deal of computer memory and storage will be spent on building and saving a SLAM map. This may not be necessary from a navigation standpoint, when only the observer pose is truly critical. It may be advisable to de-couple localization and mapping, and to first solve one problem, then the other. Generation of a map is certainly a useful scientific objective, but does not need to be part of the critical path when executing a traverse. It is the opinion of this author that such methods may be better suited to after-the-fact analysis of images or video sent down after a mission or traverse has been executed.

HAZARD DETECTION AND PATH-PLANNING

Hazard Detection and Avoidance is a critical capability that has flown on every Mars rover[17,22,27,31] and will be included on rovers still under development.[39] Hazards include rocks (or anything the rover could bump into), slopes (or anything that could cause the rover to tip), and loose terrain (or anything that could cause the rover wheels to slip). The most common strategy is to take a stereo pair of images with calibrated cameras having a known baseline (inter-camera distance). Points between these images can be correlated, and their corresponding 3D points triangulated, resulting in a 3D map of the local environment (typically a point cloud or a DEM). One could also arrive at a point cloud using a LIDAR system, though this has not yet been done in practice. These data can be processed to detect hazardous rocks, slopes, or even craters—which are useful for navigation but also pose slope hazards themselves.

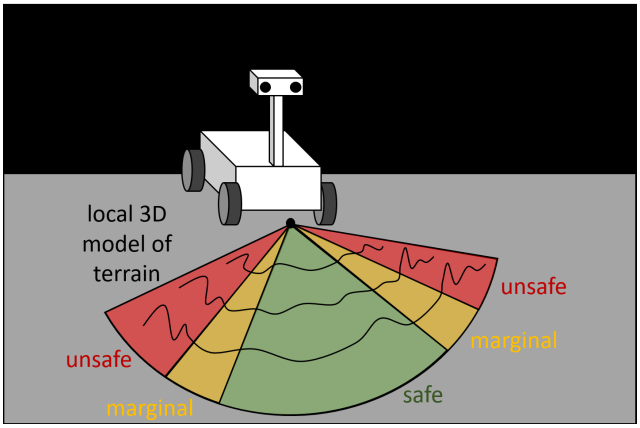

Figure 7. A point cloud or DEM of the local environment contains enough information to assess which path forward is most safe.

Early versions of rover path planning consisted of fitting a plane the size of the total rover footprint onto the local hazard map, looking for safe regions of relatively flat terrain.[21] This results in evaluation of immediate paths forward rated by traversal speed and ease, as is illustrated in Fig. 7. Recent improvements use a more complex rover shape model, making analysis more computationally expensive but allowing rovers to traverse riskier environments.[31,60] Historically, image processing required rovers to stop to for a period of time before proceeding.[61] The Perseverance rover uses a dedicated FPGA to speed up processing, allowing Percy to move at its full speed of just over 4 cm/s while performing HDA in real-time.[31] Perseverance's hazard detection is so reliable that the rover has been trusted to navigate across terrain that no human eye has ever seen, finding its own safe way there. Inevitably, a rover will encounter previously unseen hazards along its planned traverse. Hazard detection is therefore a continuous and ongoing process..

The most subtle hazard is loose terrain that would allow a rover to "spin its wheels" and ultimately become stuck. This problem was encountered on the MER mission when Opportunity was commanded to drive 50 meters but became stuck on "Purgatory Ripple." Visual odometry was used to determine that the rover had actually only moved 2 meters.[22,23] Thankfully, Oppy was able to back up and work its way around this loose terrain feature. It has been noted that slip hazards are sometimes related to patches of otherwise featureless terrain. Other times, a rock comes loose and moves under the rover as it is commanded to drive forward.[25] The detection of terrain which might pose a slip hazard is still an active area of research and development.

STELLAR POSITIONING SYSTEM

A modern adaptation of nautical celestial navigation (i.e., using a sextant and star charts) and can be accomplished with well-understood aerospace navigation hardware. Celestial navigation was discussed in the context of the Apollo missions[9,10] but was never implemented in favor of simplicity. In recent years, this method has come to be known as the Stellar Positioning System (SPS)[62,63] and has undergone promising hardware tests.[64] This method of navigation requires only a star tracker and an IMU, as illustrated in Fig. 8 and could return a position estimate in a lost-on-the-surface scenario.

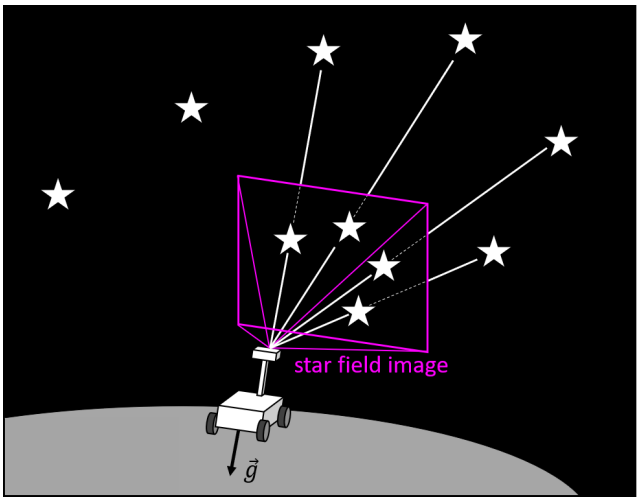

Figure 8. If an observer measures their attitude with respect to inertial space (e.g., using a star tracker) and their local apparent "down" direction, they have enough information to deduce their global position.

The relationship between the inertial (world) frame and the Moon-Centered, Moon-Fixed Frame (MCMF) is simply a function of time. The rotation between the inertial frame and the body frame of an observer can be estimated using a star tracker. The local gravity vector can be measured in the body frame with an IMU. A few frame transformations yield the local gravity vector in the MCMF frame, which contains enough information to estimate latitude and longitude. An accurate model of the Moon's gravity is necessary to realize this method. The local gravity vector would only point opposite an observer's position vector on a celestial body with a perfectly spherical gravitational field. A navigator must account for the "lumpiness" of the Moon's gravitational field using an empirically derived gravity model. In addition, the interlock angle between the vehicle's star tracker and IMU must be known with great precision. Even if this were precisely calibrated on the ground before launch, it would almost certainly need to be re-calibrated multiple times throughout the mission due to vibration and thermal effects.

Recently, a hardware demonstration of SPS was performed at NASA MSFC using a commercially available accelerometer, star tracker, and clock.[64] They started by solving for a rough estimate of their position using a spherically symmetric gravity model, then began iterating using higher-fidelity gravity models until their position converged. Using measurements taken over the course of an hour, the team estimated their position with $\approx$ 421 meters error. Amert et. al.[64] are optimistic that a similar setup on the Moon could achieve $\approx$ 125 meters position error (scaling by the relative size of the Earth and the Moon).

SURFACE CRATER-BASED NAVIGATION

Crater-based navigation for a spacecraft in orbit has been well-studied and is still an area of research and development.[65–67] Crater-based navigation can generally be divided into four parts: (1) crater detection from available data, (2) building a reference catalog of craters, (3) identifying craters (i.e., matching observed and catalog craters), and (4) estimating observer pose (e.g., solving the perspective-n-point problem[53,68]). Crater detection can be done using either neural networks[69,70] or classical image processing techniques.[71] A wealth of crater detection algorithms already exist.[72–74] Crater maps, such as the Robbins Crater Database,[75] are typically generated for scientific purposes but are perfectly useful to a navigator. Crater identification can be reduced to a database search problem through the use of crater invariants—numerical nametags given to groups of craters which are the same if computed from a catalog or from information in an image.[65]

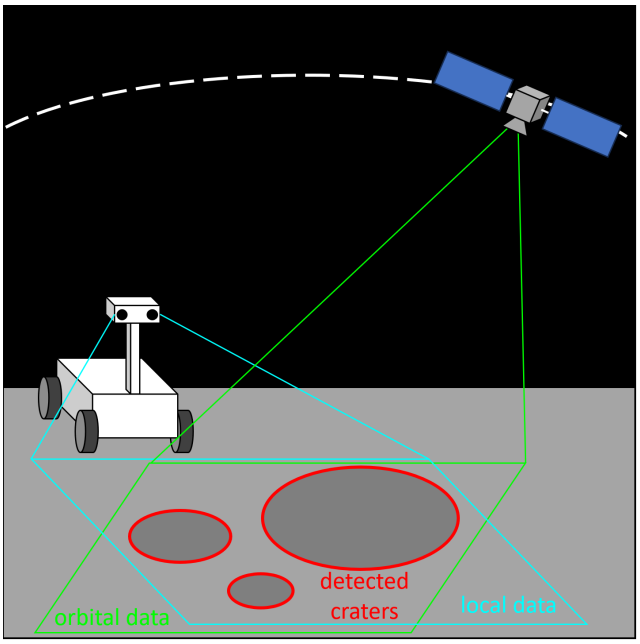

Figure 9. Craters seen by an observer on the surface can be correlated with craters seen from orbit, and this can help estimate an observer's pose.

Recent work at JPL has looked into extending this strategy to lunar surface operations,[76–78] although this poses some unique challenges. The algorithm is known as LunarNav[77] and the general premise is illustrated in Fig. 9. Craters must first be detected in mono images or in point clouds generated from a stereo camera pair or LIDAR instrument—two problems which do not have mature or well-tested solutions. The area around an observer may not feature enough craters for a pose solution, and craters that can be seen from the surface are generally too small to be included in a global crater database. Thus, navigators must construct a mission-specific crater database using orbital images. The invariant-based crater ID method described by Christian[65] relies on the assumption that the craters are either co-planar or are all tangent to the surface of a tri-axial ellipsoid. These assumptions do not necessarily hold on the surface. Thus, the authors of LunarNav are forced to use an ad-hoc technique for crater identification. LunarNav and other methods of lunar surface navigation using craters are very promising and are certainly worth further research and development, but are not as mature as some of the other surface navigation technologies discussed in this paper.

SURFACE SKYLINE-BASED NAVIGATION

Features in the silhouette formed by the distant terrain against the background of space can be used for localization. This has been studied in the context of lunar[79] and Martian[80] exploration, and is being actively researched at NASA JSC and Goddard Space Flight Center (GSFC).[81,82] Skyline-based navigation can be broken down into two basic strategies: coarse pose estimation by observing the entire skyline, and fine pose estimation looking at particular features in an image.

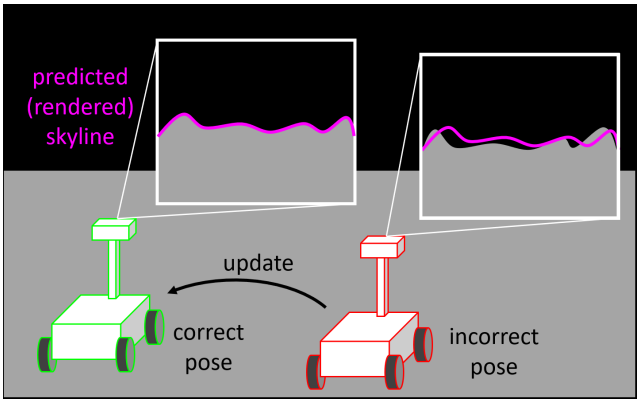

Figure 10. The skyline acts as a fingerprint—it is unique to a particular observer position. Thus, a lookup table of horizon lines can facilitate localization in a lost-on-the-surface scenario.

The panoramic skyline (or partial skyline) does not contain enough information for observer localization. The skyline is more like a fingerprint—unique to one location, but only useful if there is a database of other skylines to compare to. The (pre-rendered) skyline one would expect to see matches observations only at one location on the surface, as is illustrated in Fig. 10. Coarse pose estimation using horizon images can be broken down into four steps. The first is using DEM data to render skylines as they would be seen at several reference points on the surface. The second is image processing to find horizon points, a task made more difficult by the extreme lighting conditions at the lunar south pole. The third is combining images to form a panorama, while being careful to avoid introducing any image stitching artifacts. Finally, the fourth is comparing the observed skyline to many rendered skylines and returning the best match. The accuracy of this coarse pose estimation technique is limited by the number of reference points considered in the first step. One is forced to assume they are at, or near, one of these points. There will always be some nontrivial pose error, and the possibility of a false-positive match is non-negligible. Still, this gives the observer a coarse pose estimate even when lost on the surface of another celestial body.

A second form of skyline-based navigation can be used to refine an initial pose estimate. An observer can render the expected skyline corresponding to a specific pose, then correlate features between the expected and measured skylines. This correlation problem is nontrivial, and remains an area of active research.[82,83] By contrast, correlation between rendered image points and 3D points taken from a reference DEM is trivial. The result of this effort is a set of direction vectors to known 3D points, and one only needs to solve the PnP problem to estimate observer pose. This process can be repeated as necessary until pose converges. Coarse and fine skyline-based navigation dovetail nicely. Coarse skyline-based navigation can generate a rough solution even when lost on the surface. Fine skyline-based navigation can dramatically improve an initial pose estimate. These techniques are promising, but have only been studied in theoretical and academic contexts.

SURFACE NAVIGATION USING ORBITAL ELEVATION MAPS

Hazard detection has been done on Mars rovers[17,27,31] and Moon rovers[84] using stereo cameras and on terrestrial rovers using LIDAR.[85] The use of LIDAR for HDA was also explored on the Morpheus lander[86]—a part of the Autonomous Landing and Hazard Avoidance (ALHAT) project.[87] While stereo cameras have greater range than LIDAR under good lighting conditions, LIDAR works under all lighting conditions,[88,89] making it particularly attractive for exploring the lunar south pole. At present, no rover has flown to another celestial body equipped with a LIDAR, but NASA Goddard is working on a Space Qualified Rover LIDAR (SQRLi). A point cloud or DEM of local terrain can be used for navigation as well as HDA. A measured local DEM should fit onto a DEM generated from orbital imagery like a puzzle piece, allowing for observer pose estimation, as illustrated in Fig. 11. Global pose estimation from LIDAR data has been studied for landers[89] and has been demonstrated on Earth with helicopter flight test data.[88] This has also been investigated for maritime navigation—a ship can navigate using 3D scans of the ocean floor.[90]

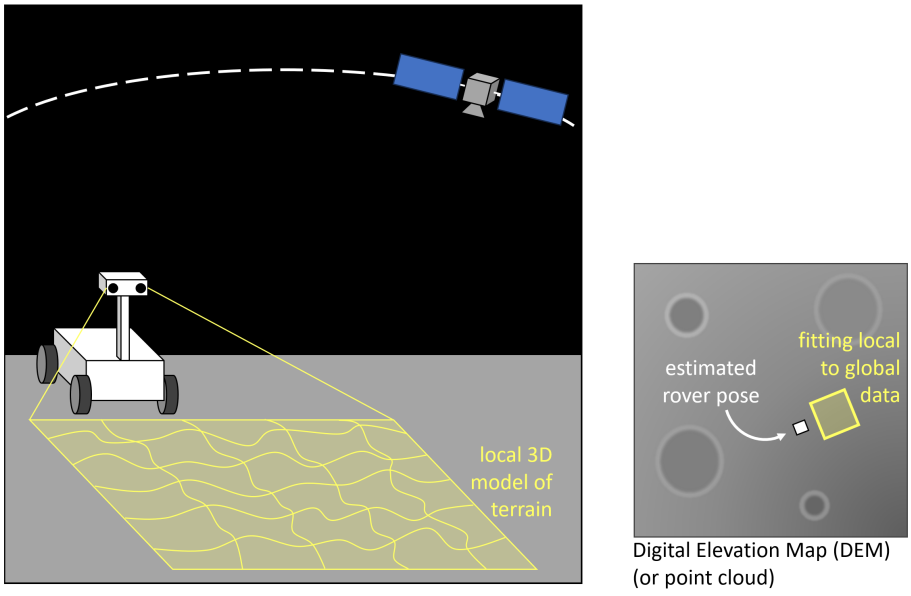

Figure 11. A stereo image pair or LIDAR scan can be used to create a point cloud or DEM of the local environment. This is useful for hazard detection and also for navigation, by matching the local DEM to a global DEM made from orbital images.

The most common strategy involves starting with an initial guess of observer position and iterating until measured and reference DEMs align perfectly, using something like a floating point correction algorithm.[88] In the absence of an initial guess, the DEM correspondence problem becomes much harder to solve. The naïve solution would be a brute force search of possible observer poses to align DEMs. The Standalone Hazard Evaluation by Refinement of Instrument Findings (SHERIF) algorithm[91] is being developed at NASA JSC, and is also focused on point cloud analysis for HDA. One of the key advancements in the SHERIF algorithm is the development of 3D feature points that can be used to stitch together LIDAR point clouds. These are essentially the 3D version of SIFT or other 2D image features. These features could be used to match a small (measured) DEM to a large (reference) DEM in a fraction of the time of a brute force search. While promising, this method of navigation has not yet been implemented and remains low-TRL.

SURFACE NAVIGATION USING ORBITAL IMAGERY

Until very recently, autonomous global localization for rovers has not been a consideration. The only missions that ventured far enough from their landing site to require global localization have been Mars rovers, and those have historically performed global localization by sending images to Earth for processing. Engineers at JPL have found a way to automate this process using a modified form of the census transform. Their algorithm, named "Censible" has been run on the Perseverance rover flight computer—the first demonstration of on-board global localization on another world.[92,93]

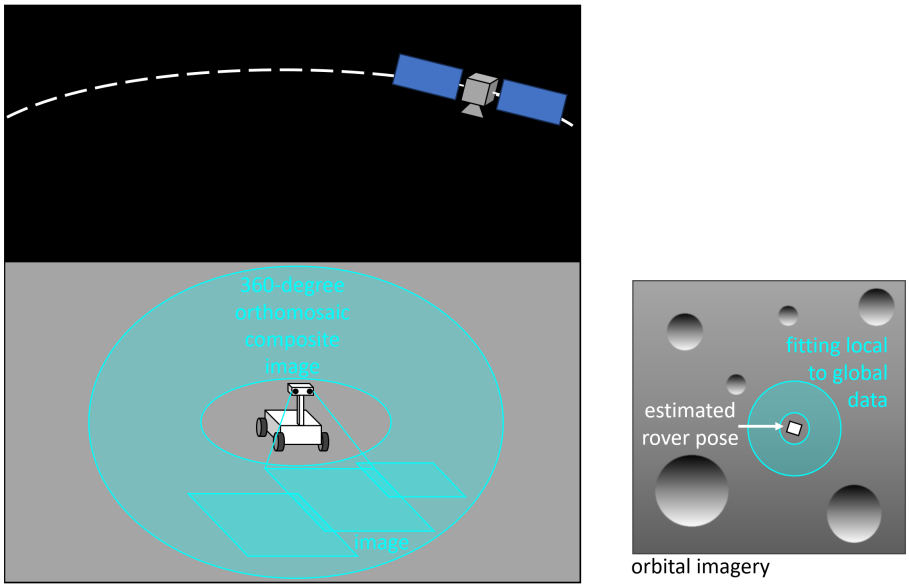

Figure 12. The Censible algorithm involves producing a synthetic top-down image of the local terrain, then comparing this to orbital imagery for observer pose estimation.

A stereo camera pair can take an ensemble of stereo images in a panorama around the observer, and these can be used to synthesize a 3D point cloud around the vehicle. Color rectification must be performed on each image, then a brightness can be assigned to each point in the point cloud. This results in an orthomosaic image which should fit like a puzzle piece onto a larger orbital image of the surrounding region. A comparison is then performed between the synthetic local and orbital images of the surface using a modified census transform. The brightness at each pixel is compared to its surrounding eight neighbors, and a descriptor is built for each pixel. This comparison is far more robust than the normalized cross correlation method. The search space can also be reduced by limiting the region of the orbital image under consideration to a spot near some initial guess of the observer position. Once the center of the orthomosaic is correlated to a location in an orbital image it is straightforward to estimate the latitude and longitude of the observer.

Much like DEM-based navigation and skyline-based navigation, this method relies on a comparison of local to global data after some modification that makes a direct "apples-to-apples" comparison possible. The Censible algorithm can be carried out by any rover with a stereo camera pair and a mast that allows for full 360-degree rotation. It requires a reasonable amount of on-board processing power (certainly less than is required for real-time hazard avoidance). This technique exists only as a flight demonstration, but it is sensible to assume this, or something similar, will become much more prevalent on future missions.

SUMMARY AND CONCLUSION

Case studies of the Apollo LRV, the family of Mars rovers, and the (still in-development) VIPER rover provide insight into the challenges of surface navigation. This paper reviewed methods of autonomous global navigation (finding latitude and longitude), local navigation (finding distance and direction traveled), and HDA. These are summarized in Tab. 1, along with their flight heritage and required hardware. It is the hope of this author that the information in this paper will prove useful for those involved in the design, development, and testing of future rovers.

Table 1. Summary of autonomous surface navigation techniques discussed in this paper.

Navigation Technique	Hardware Required	Considerations
Inertial Navigation and Wheel Odometry (local navigation)	wheel odometers, IMU, suspension encoders (opt.),	wheel slip, gyro drift
Visual Odometry (local navigation)	stereo camera pair or mono camera with, range instrument FPGA or GPU (opt.)	choice of feature points point correspondence
SLAM (local navigation)	mono/stereo cameras or LIDAR, IMU (opt.), FPGA or GPU (opt.)	choice of feature points point correspondence requires loop closure computer storage limitations (no space flight heritage)
HDA and Path Planning	stereo camera pair or LIDAR, IMU, FPGA or GPU (opt.)	computationally expensive necessary throughout traverse
Celestial Navigation (e.g., SPS) (global navigation)	star tracker, IMU, clock	must know alignment from IMU to star tracker (no space flight heritage)
Surface Crater Nav (e.g., LunarNav) (global navigation)	mono/stereo cameras or LIDAR, FPGA or GPU (opt.)	surface crater ID problem, relies on dense crater distribution (no space flight heritage)
Skyline Navigation (global navigation)	mono/stereo cameras, FPGA or GPU (opt.)	must pre-render skylines or be able to render onboard, point correspondence (no space flight heritage)
DEM Matching (e.g., SHERIF) (global navigation)	stereo cameras or LIDAR, FPGA or GPU (opt.)	point correspondence (no space flight heritage)
Orbital Image Matching (e.g., Censible) (global navigation)	stereo cameras FPGA or GPU (opt.)	computationally expensive (minimal space flight heritage)

REFERENCES

[1] C. L. Thornton and J. S. Border, *Radiometric Tracking Techniques for Deep-Space Navigation*. New York City, NY, USA: John Wiley & Sons, 2003, 10.1002/0471728454.

[2] L. B. Winternitz, W. A. Bamford, A. C. Long, and M. Hassouneh, "GPS based autonomous navigation study for the lunar gateway," *Annual American Astronautical Society (AAS) Guidance, Navigation, and Control Conference*, No. AAS 19-096, 2019.

[3] A. Delépaut, P. Giordano, J. Ventura-Traveset, D. Blonski, M. Schönfeldt, P. Schoonejans, S. Aziz, and R. Walker, "Use of GNSS for lunar missions and plans for lunar in-orbit development," *Advances in Space Research*, Vol. 66, No. 12, 2020, pp. 2739–2756, 10.1016/j.asr.2020.05.018.

[4] J. Esper, "Draft LunaNet Interoperability Specification," Tech. Rep. TP-20210021073/Rev.2, National Aeronautics and Space Administration, Mar. 2022.

[5] E. Anzalone, A. Iyer, and T. Statham, "Use of navigation beacons to support lunar vehicle operations," *2020 IEEE Aerospace Conference*, IEEE, 2020, pp. 1–13, 10.1109/AERO47225.2020.9172663.

[6] D. Christensen and D. Geller, "Terrain-relative and beacon-relative navigation for lunar powered descent and landing," *The Journal of the Astronautical Sciences*, Vol. 58, No. 1, 2011, pp. 121–151, 10.1007/BF03321162.

[7] H. W. Tindall, "Rover Navigation," Tech. Rep. 69-PA-T-147A, NASA Manned Spacecraft Center, Nov. 1969.

[8] W. G. Heffron and F. LaPiana, "The Navigation System of the Lunar Roving Vehicle," Tech. Rep. TM-70-2014-8, Bellcomm, Inc., Dec. 1970.

[9] M. Streicher, "Lunar Surface Navigation for a Roving Vehicle (MOLAB)," *The Space Congress Proceedings*, No. 1, 1965.

[10] T. T. Trexler and R. B. Odden, "Lunar Surface Navigation," *IEEE Transactions on Aerospace and Electronic Systems*, No. 6, 1966, pp. 252–259, 10.1109/TAES.1966.4502015.

[11] I. M. Salzberg, "Tracking the Apollo Lunar Rover with Interferometry Techniques," *Proceedings of the IEEE*, Vol. 61, No. 9, 1973, pp. 1233–1236, 10.1109/PROC.1973.9251.

[12] E. Swift, *Across the Airless Wilds*. Worthington, OH: Custom House, 2021.

[13] I. Y. Bar-Itzhack and M. L. Carothers, "LRV Navigation Support at Bellcomm During the Apollo 15 Mission," Tech. Rep. B71 08026, Bellcomm, Inc., Aug. 1971.

[14] E. C. Smith and W. C. Mastin, "Lunar roving vehicle navigation system performance review," Tech. Rep. TN D-7469, NASA Marshall Space Flight Center, Nov. 1973.

[15] A. H. Mishkin, J. C. Morrison, T. T. Nguyen, H. W. Stone, B. K. Cooper, and B. H. Wilcox, "Experiences with operations and autonomy of the mars pathfinder microrover," *1998 IEEE aerospace conference proceedings (Cat. No. 98TH8339)*, Vol. 2, IEEE, 1998, pp. 337–351, 10.1109/AERO.1998.687920.

[16] B. Wilcox and T. Nguyen, "Sojourner on mars and lessons learned for future planetary rovers," tech. rep., SAE Technical Paper, 1998, 10.4271/981695.

[17] M. Bajracharya, M. W. Maimone, and D. Helmick, "Autonomy for mars rovers: Past, present, and future," *Computer*, Vol. 41, No. 12, 2008, pp. 44–50, 10.1109/MC.2008.479.

[18] R. Team, "The Pathfinder microrover," *Journal of Geophysical Research: Planets*, Vol. 102, No. E2, 1997, pp. 3989–4001, 10.1029/96JE01922.

[19] J. Matijevic, "Autonomous navigation and the sojourner microrover," *Science*, Vol. 280, No. 5362, 1998, pp. 454–455, 10.1126/science.280.5362.454.

[20] J. Maki, J. Bell III, K. E. Herkenhoff, S. Squyres, A. Kiely, M. Klimesh, M. Schwochert, T. Litwin, R. Willson, A. Johnson, *et al.*, "Mars exploration rover engineering cameras," *Journal of Geophysical Research: Planets*, Vol. 108, No. E12, 2003, 10.1029/2003JE002077.

[21] M. Maimone, J. Biesiadecki, E. Tunstel, Y. Cheng, and C. Leger, "Surface navigation and mobility intelligence on the Mars Exploration Rovers," *Intelligence for space robotics*, 2006, pp. 45–69.

[22] J. J. Biesiadecki and M. W. Maimone, "The mars exploration rover surface mobility flight software: driving ambition," *2006 IEEE Aerospace Conference*, IEEE, 2006, pp. 15–pp, 10.1109/AERO.2006.1655723.

[23] M. Maimone, Y. Cheng, and L. Matthies, "Two years of visual odometry on the mars exploration rovers," *Journal of Field Robotics*, Vol. 24, No. 3, 2007, pp. 169–186, 10.1002/rob.20184.

[24] M. Maimone, A. Johnson, Y. Cheng, R. Willson, and L. Matthies, "Autonomous navigation results from the Mars Exploration Rover (MER) mission," *Experimental Robotics IX: The 9th International Symposium on Experimental Robotics*, Springer, 2006, pp. 3–13, 10.1007/11552246_1.

[25] A. Rankin, M. Maimone, J. Biesiadecki, N. Patel, D. Levine, and O. Toupet, "Driving curiosity: Mars rover mobility trends during the first seven years," *2020 IEEE Aerospace Conference*, IEEE, 2020, pp. 1–19, 10.1109/AERO47225.2020.9172469.

[26] J. Maki, D. Thiessen, A. Pourangi, P. Kobzeff, T. Litwin, L. Scherr, S. Elliott, A. Dingizian, and M. Maimone, "The Mars science laboratory engineering cameras," *Space science reviews*, Vol. 170, 2012, pp. 77–93, 10.1007/s11214-012-9882-4.

[27] J. P. Grotzinger, J. Crisp, A. R. Vasavada, R. C. Anderson, C. J. Baker, R. Barry, D. F. Blake, P. Conrad, K. S. Edgett, B. Ferdowski, *et al.*, "Mars Science Laboratory mission and science investigation," *Space science reviews*, Vol. 170, 2012, pp. 5–56, 10.1007/s11214-012-9892-2.

[28] N. Warner, M. Silverman, J. Samuels, L. DeFlores, A. Sengstacken, J. Maki, A. Scodary, S. Peters, T. Litwin, and B. Metz, "The Mars Science Laboratory Remote Sensing Mast," *2016 IEEE Aerospace Conference*, IEEE, 2016, pp. 1–9, 10.1109/AERO.2016.7500554.

[29] K. A. Farley, K. H. Williford, K. M. Stack, R. Bhartia, A. Chen, M. d. l. Torre, K. Hand, Y. Goreva, C. D. Herd, R. Hueso, *et al.*, "Mars 2020 mission overview," *Space Science Reviews*, Vol. 216, 2020, pp. 1–41, 10.1007/s11214-020-00762-y.

[30] J. Maki, D. Gruel, C. McKinney, M. Ravine, M. Morales, D. Lee, R. Willson, D. Copley-Woods, M. Valvo, T. Goodsall, *et al.*, "The Mars 2020 Engineering Cameras and microphone on the perseverance rover: A next-generation imaging system for Mars exploration," *Space science reviews*, Vol. 216, 2020, pp. 1–48, 10.1007/s11214-020-00765-9.

[31] V. Verma, M. W. Maimone, D. M. Gaines, R. Francis, T. A. Estlin, S. R. Kuhn, G. R. Rabideau, S. A. Chien, M. M. McHenry, E. J. Graser, *et al.*, "Autonomous robotics is driving Perseverance rover's progress on Mars," *Science Robotics*, Vol. 8, No. 80, 2023, 10.1126/scirobotics.adi3099.

[32] R. Rieber, M. McHenry, P. Twu, and M. M. Stragier, "Planning for a Martian road trip-the Mars2020 mobility systems design," *2022 ieee aerospace conference (aero)*, IEEE, 2022, pp. 01–18, 10.1109/AERO53065.2022.9843375.

[33] J. Balaram, M. Aung, and M. P. Golombek, "The ingenuity helicopter on the perseverance rover," *Space Science Reviews*, Vol. 217, No. 4, 2021, p. 56, 10.1007/s11214-021-00815-w.

[34] D. S. Bayard, D. T. Conway, R. Brockers, J. H. Delaune, L. H. Matthies, H. F. Grip, G. B. Merewether, T. L. Brown, and A. M. San Martin, "Vision-based navigation for the NASA mars helicopter," *AIAA Scitech 2019 Forum*, 2019, p. 1411, 10.2514/6.2019-1411.

[35] H. F. Grip, D. Conway, J. Lam, N. Williams, M. P. Golombek, R. Brockers, M. Mischna, and M. R. Cacan, "Flying a helicopter on Mars: How ingenuity's flights were planned, executed, and analyzed," *2022 IEEE Aerospace Conference (AERO)*, IEEE, 2022, pp. 1–17, 10.1109/AERO53065.2022.9843813.

[36] J. A. Christian, L. Hong, P. McKee, R. Christensen, and T. P. Crain, "Image-based lunar terrain relative navigation without a map: Measurements," *Journal of Spacecraft and Rockets*, Vol. 58, No. 1, 2021, pp. 164–181, 10.2514/1.A34875.

[37] A. Colaprete, D. Andrews, W. Bluethmann, R. C. Elphic, B. Bussey, J. Trimble, K. Zacny, and J. E. Captain, "An overview of the volatiles investigating polar exploration rover (viper) mission," *AGU fall meeting abstracts*, Vol. 2019, 2019, pp. P34B–03.

[38] D. Andrews, "VIPER: Systems Integration Status," *74th International Astronautical Congress (IAC)*, Baku, AZ, Oct. 2023.

[39] A. Colaprete, "Volatiles Investigating Polar Exploration Rover (VIPER)," Tech. Rep. 20210015009, NASA Ames Research Center, 2021.

[40] R. A. Beyer, U. Wong, L. Edwards, A. Colaprete, K. Ennico-Smith, K. Jeppesen, J. Joy, M. O'Connor, R. Reed, B. Wright, P. Becerra, M. Nilsson, G. Vives, S. Benekos, , and S. Beauvivre, "VIPER Visible Imaging System," *53rd Lunar and Planetary Science Conference*, Lunar Planetary Institute, 2022, p. 2466.

[41] T. Fong, "Volatiles Investigating Polar Exploration Rover (VIPER)," *West Virginia University Robotics Seminar*, Nov. 2021.

[42] R. Vaughan, "VIPER–Volatiles Investigating Polar Exploration Rover: Mission Overview," *International Small Satellite Conference*, Virtual, May 2020.

[43] M. Fazekas, P. Gáspár, and B. Németh, "Calibration and improvement of an odometry model with dynamic wheel and lateral dynamics integration," *Sensors*, Vol. 21, No. 2, 2021, p. 337, 10.3390/s21020337.

[44] C. Kilic, J. N. Gross, N. Ohi, R. Watson, J. Strader, T. Swiger, S. Harper, and Y. Gu, "Improved planetary rover inertial navigation and wheel odometry performance through periodic use of zero-type constraints," *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2019, pp. 552–559, 10.1109/IROS40897.2019.8967634.

[45] M. Brossard and S. Bonnabel, "Learning wheel odometry and IMU errors for localization," *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, 2019, pp. 291–297, 10.1109/ICRA.2019.8794237.

[46] F. Rogers-Marcovitz, M. George, N. Seegmiller, and A. Kelly, "Aiding off-road inertial navigation with high performance models of wheel slip," *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2012, pp. 215–222, 10.1109/IROS.2012.6385701.

[47] D. G. Lowe, "Object recognition from local scale-invariant features," *Proceedings of the seventh IEEE international conference on computer vision*, Vol. 2, IEEE, 1999, pp. 1150–1157, 10.1109/ICCV.1999.790410.

[48] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, Vol. 60, 2004, pp. 91–110, 10.1023/B:VISI.0000029664.99615.94.

[49] H.-J. Chien, C.-C. Chuang, C.-Y. Chen, and R. Klette, "When to use what feature? SIFT, SURF, ORB, or A-KAZE features for monocular visual odometry," *2016 International Conference on Image and Vision Computing New Zealand (IVCNZ)*, IEEE, 2016, pp. 1–6, 10.1109/IVCNZ.2016.7804434.

[50] S. A. K. Tareen and Z. Saleem, "A comparative analysis of sift, surf, kaze, akaze, orb, and brisk," *2018 International conference on computing, mathematics and engineering technologies (iCoMET)*, IEEE, 2018, pp. 1–10, 10.1109/ICOMET.2018.8346440.

[51] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, Vol. 24, No. 6, 1981, pp. 381–395, 10.1145/358669.358692.

[52] S. Andolfo, F. Petricca, and A. Genova, "Visual Odometry analysis of the NASA Mars 2020 Perseverance rover's images," *2022 IEEE 9th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, IEEE, 2022, pp. 287–292, 10.1109/MetroAeroSpace54187.2022.9856188.

[53] T. M. Howard, A. Morfopoulos, J. Morrison, Y. Kuwata, C. Villalpando, L. Matthies, and M. McHenry, "Enabling continuous planetary rover navigation through FPGA stereo and visual odometry," *2012 IEEE Aerospace Conference*, IEEE, 2012, pp. 1–9, 10.1109/AERO.2012.6187041.

[54] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part I," *IEEE robotics & automation magazine*, Vol. 13, No. 2, 2006, pp. 99–110, 10.1109/MRA.2006.1638022.

[55] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *IEEE robotics & automation magazine*, Vol. 13, No. 3, 2006, pp. 108–117, 10.1109/MRA.2006.1678144.

[56] T. Taketomi, H. Uchiyama, and S. Ikeda, "Visual SLAM algorithms: A survey from 2010 to 2016," *IPSJ transactions on computer vision and applications*, Vol. 9, 2017, pp. 1–11, 10.1186/s41074-017-0027-2.

[57] A. Macario Barros, M. Michel, Y. Moline, G. Corre, and F. Carrel, "A comprehensive survey of visual slam algorithms," *Robotics*, Vol. 11, No. 1, 2022, p. 24, 10.3390/robotics11010024.

[58] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.

[59] S. Ullman, "The interpretation of structure from motion," *Proceedings of the Royal Society of London. Series B. Biological Sciences*, Vol. 203, No. 1153, 1979, pp. 405–426, 10.1098/rspb.1979.0006.

[60] K. Otsu, G. Matheron, S. Ghosh, O. Toupet, and M. Ono, "Fast approximate clearance evaluation for rovers with articulated suspension systems," *Journal of Field Robotics*, Vol. 37, No. 5, 2020, pp. 768–785, 10.1002/rob.21892.

[61] R. Volpe, T. Estlin, S. Laubach, C. Olson, and J. Balaram, "Enhanced mars rover navigation techniques," *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, Vol. 1, IEEE, 2000, pp. 926–931, 10.1109/ROBOT.2000.844167.

[62] J. J. Parish, A. S. Parish, M. Swanzy, D. Woodbury, D. Mortari, and J. L. Junkins, "Stellar positioning system (part I): an autonomous position determination solution," *Navigation*, Vol. 57, No. 1, 2010, pp. 1–12, 10.1002/j.2161-4296.2010.tb01763.x.

[63] D. P. Woodbury, J. J. Parish, A. S. Parish, M. Swanzy, R. Denton, D. Mortari, and J. L. Junkins, "Stellar positioning system (Part II): Improving accuracy during implementation," *Navigation*, Vol. 57, No. 1, 2010, pp. 13–24, 10.1002/j.2161-4296.2010.tb01764.x.

[64] J. Amert and M. Fritzinger, "Hardware Demonstration and Improvements of the Stellar Positioning System," *45th Annual AAS Guidance, Navigation and Control (GN&C) Conference*, Breckenridge, CO, Feb. 2023.

[65] J. A. Christian, H. Derksen, and R. Watkins, "Lunar crater identification in digital images," *The Journal of the Astronautical Sciences*, Vol. 68, No. 4, 2021, pp. 1056–1144, 10.1007/s40295-021-00287-8.

[66] F. C. Hanak, *Lost in low lunar orbit crater pattern detection and identification*. PhD thesis, University of Texas at Austin, 2009.

[67] L. Singh and S. Lim, "On lunar on-orbit vision-based navigation: Terrain mapping, feature tracking driven EKF," *AIAA Guidance, Navigation and Control Conference and Exhibit*, No. AIAA 2008-6834, Honolulu, HI, Aug. 2008, 10.2514/6.2008-6834.

[68] V. Lepetit, F. Moreno-Noguer, and P. Fua, "EP n P: An accurate O (n) solution to the P n P problem," *International journal of computer vision*, Vol. 81, 2009, pp. 155–166, 10.1007/s11263-008-0152-6.

[69] A. Silburt, M. Ali-Dib, C. Zhu, A. Jackson, D. Valencia, Y. Kissin, D. Tamayo, and K. Menou, "Lunar crater identification via deep learning," *Icarus*, Vol. 317, 2019, pp. 27–38, 10.1016/j.icarus.2018.06.022.

[70] L. Downes, T. J. Steiner, and J. P. How, "Deep learning crater detection for lunar terrain relative navigation," *AIAA SciTech 2020 Forum*, 2020, p. 1838, 10.2514/6.2020-1838.

[71] Y. Cheng, A. E. Johnson, L. H. Matthies, and C. F. Olson, "Optical landmark detection for spacecraft navigation," *13th AAS/AIAA Space Flight Mechanics Meeting*, No. AAS 02-224, Jan. 2003.

[72] G. Salamunićcar and S. Lončarić, "Open framework for objective evaluation of crater detection algorithms with first test-field subsystem based on MOLA data," *Advances in Space Research*, Vol. 42, No. 1, 2008, pp. 6–19, 10.1016/j.asr.2007.04.028.

[73] D. M. DeLatte, S. T. Crites, N. Guttenberg, and T. Yairi, "Automated crater detection algorithms from a machine learning perspective in the convolutional neural network era," *Advances in Space Research*, Vol. 64, No. 8, 2019, pp. 1615–1628, 10.1016/j.asr.2019.07.017.

[74] S. Woicke, A. Moreno Gonzalez, I. El-Hajj, J. Mes, M. Henkel, and R. Klavers, "Comparison of crater-detection algorithms for terrain-relative navigation," *2018 aiaa guidance, navigation, and control conference*, 2018, p. 1601, 10.2514/6.2018-1601.

[75] S. J. Robbins and B. M. Hynek, "A new global database of Mars impact craters $\geq$ 1 km: 1. Database creation, properties, and parameters," *Journal of Geophysical Research: Planets*, Vol. 117, No. E5, 2012, 10.1029/2011JE003966.

[76] L. Matthies, S. Daftry, S. Tepsuporn, Y. Cheng, D. Atha, R. M. Swan, S. Ravichandar, and M. Ono, "Lunar rover localization using craters as landmarks," *2022 IEEE Aerospace Conference (AERO)*, IEEE, 2022, pp. 1–17, 10.1109/AERO53065.2022.9843714.

[77] S. Daftry, Z. Chen, Y. Cheng, S. Tepsuporn, S. Khattak, L. Matthies, B. Coltin, U. Naal, L. M. Ma, and M. Deans, "LunarNav: Crater-based Localization for Long-range Autonomous Lunar Rover Navigation," *2023 IEEE Aerospace Conference*, IEEE, 2023, pp. 1–15, 10.1109/AERO55745.2023.10115640.

[78] A. Cauligi, R. M. Swan, H. Ono, S. Daftry, J. Elliott, L. Matthies, and D. Atha, "Shadownav: Crater-based localization for nighttime and permanently shadowed region lunar navigation," *2023 IEEE Aerospace Conference*, IEEE, 2023, pp. 1–12, 10.1109/AERO55745.2023.10115745.

[79] E. E. Palmer, J. N. Head, R. W. Gaskell, M. V. Sykes, and B. McComas, "Mercator—Independent rover localization using stereophotoclinometry and panoramic images," *Earth and Space Science*, Vol. 3, No. 12, 2016, pp. 488–509, 10.1002/2016EA000189.

[80] S. Chiodini, M. Pertile, S. Debei, L. Bramante, E. Ferrentino, A. G. Villa, I. Musso, and M. Barrera, "Mars rovers localization by matching local horizon to surface digital elevation models," *2017 IEEE International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, IEEE, 2017, pp. 374–379, 10.1109/MetroAeroSpace.2017.7999600.

[81] A. Liounis, A. Yew, and D. McGann, "Science Geotagging and Localization Using Lunar Horizons," *2023 NASA Exploration Science Forum*, 2023.

[82] W. Driessen, S. Kaki, A. Liounis, D. McGann, P. McKee, A. Tennenbaum, and A. Yew, "Monocular Horizon Navigation," *4th Space Imaging Workshop*, 2024.

[83] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone, "Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection," *IEEE Robotics and Automation Letters*, Vol. 5, No. 2, 2020, pp. 1127–1134, 10.1109/LRA.2020.2965893.

[84] W. Wan, J. Wang, K. Di, J. Li, Z. Liu, P. Man, Y. Wang, T. Yu, C. Liu, and L. Li, "Enhanced lunar topographic mapping using multiple stereo images taken by Yutu-2 rover with changing illumination conditions," *Photogrammetric Engineering & Remote Sensing*, Vol. 87, No. 8, 2021, pp. 567–576, 10.14358/PERS.87.8.567.

[85] G. Ishigami, M. Otsuki, and T. Kubota, "Lidar-based terrain mapping and navigation for a planetary exploration rover," *Proceedings of international symposium on artificial intelligence, robotics and automation in space (i-SAIRAS)*, 2012.

[86] J. B. Olansen, "Project morpheus: Lander technology development," *AIAA SPACE 2014 Conference and Exposition*, 2014, p. 4314, 10.2514/6.2014-4314.

[87] F. Amzajerdian, D. Pierrottet, L. B. Petway, G. D. Hines, V. E. Roback, and R. A. Reisse, "Lidar sensors for autonomous landing and hazard avoidance," *AIAA Space 2013 Conference and Exposition*, 2013, p. 5312, 10.2514/6.2013-5312.

[88] A. Johnson and T. Ivanov, "Analysis and testing of a lidar-based approach to terrain relative navigation for precise lunar landing," *Aiaa guidance, navigation, and control conference*, No. AIAA 2011-6578, Portland, OR, Aug. 2011, 10.2514/6.2011-6578.

[89] T. P. Setterfield, R. A. Hewitt, P.-T. Chen, A. T. Espinoza, N. Trawny, and A. Katake, "Lidar-inertial based navigation and mapping for precision landing," *2021 IEEE Aerospace Conference (50100)*, IEEE, 2021, pp. 1–19, 10.1109/AERO50100.2021.9438153.

[90] J. Olofsson, G. Hendeby, F. Gustafsson, D. Maas, and S. Maranò, "GNSS-free maritime navigation using radar and digital elevation models," *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*, Rustenburg, South Africa, IEEE, 2020, pp. 1–8, 10.23919/FUSION45008.2020.9190450.

[91] K. R. Kobylka and D. Adams, "Standalone Hazard Evaluation and Refinement From Instrument Findings (SHERIF)," *2024 AAS Rocky Mountain GNC Conference*, 2024.

[92] J. Nash, Q. Dwight, L. Saldyt, H. Wang, S. Myint, A. Ansar, and V. Verma, "Censible: A Robust and Practical Global Localization Framework for Planetary Surface Missions," *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 8642–8648, 10.1109/ICRA57147.2024.10611697.

[93] V. Verma, J. Nash, L. Saldyt, Q. Dwight, H. Wang, S. Myint, J. Biesiadecki, M. Maimone, A. Tumbar, A. Ansar, *et al.*, "Enabling Long & Precise Drives for The Perseverance Mars Rover via Onboard Global Localization," *2024 IEEE Aerospace Conference*, IEEE, 2024, pp. 1–18, 10.1109/AERO58975.2024.10521160.