Efficient Stereo Vision for Feature-sparse Lunar Images

15-300 Project Proposal

William Lee-Moore

Project Description

In this project, I will perform a study comparing a range of different algorithms for automated perception of the surrounding environment based on stereo vision. The goal of the study is to identify one "final recommendation" algorithm that best meets the needs of MoonRanger. The usefulness of this final recommendation can be defined by three criteria as follows:

- Producing results with desired format and content for the purposes of MoonRanger
- Producing results with adequate accuracy for the purposes of MoonRanger
- Running with sufficient efficiency for MoonRanger hardware

This project proposal now describes necessary background information about MoonRanger, then discusses each of these criteria in detail and how to achieve them.

MoonRanger Background Information

MoonRanger is a project intending to put an autonomous robotic rover, called MoonRanger, on the moon (Spice 2019). This is a recent project, started in late summer of 2019 (Spice 2019). MoonRanger will by much smaller and cheaper than other outer-space rovers and will rely exclusively on solar power rather than using a radioisotope heat source (MoonRanger 2019). Unlike all previous outer-space rovers with the exception of Pathfinder, MoonRanger will not carry an onboard radio capable of communication with Earth (MoonRanger 2019). Unlike Pathfinder, MoonRanger's primary mission will take it outside of communications range with its lander, meaning it will spend time operating without any communications link to Earth (MoonRanger 2019). During this time, it will have to be entirely autonomous, including sensing the environment around it automatically with stereo vision to identify safe paths that avoid obstacles.

The project is led by Professor William (Red) Whittaker. Professor David Wettergreen is also closely involved with the software side of the project owing to his extensive research into autonomy for planetary rovers (Wettergreen and Wagner 2012) (Wettergreen, Foil, et al. 2014). Senior Project Scientist Heather Jones is also closely involved with the project, especially in the role of coordinating the student teams. For my project, I will work closely with the Autonomy and Navigation Software Teams, especially the Mapping and Planning teams. The work by the Avionics team on camera and computer hardware, and by the Systems team on project timeline and risk management is also highly relevant to this project.

Criterion 1: Desired Format and Content for MoonRanger

The first criterion for the final recommendation is "producing results with desired format and content for the purposes of MoonRanger." This criterion is entirely about the theoretical idealized output of the algorithm. The other two criteria assess the effectiveness with which the algorithm transforms input data to output results; this criterion assesses whether the output results meet the needs of MoonRanger. There are a large number of ways that the output results can meet or fail to meet these needs:

Providing Sufficient Information for Obstacle Avoidance

One of the primary purposes of stereo vision in MoonRanger is to identify all obstacles automatically so that the rover can automatically plan routes that avoid obstacles. As such, one objective inside this criterion is that the output results of the final recommendation algorithm should provide enough information to enable this obstacle avoidance.

To measure success or failure in this objective, a definition of what obstacles should be identified is needed. Attempting to detect all hazards is a poor approach to take. For example, the Mars Exploration Rovers Spirit and Opportunity used computer vision to automatically identify geometric hazards to the rovers (such as large rocks and dangerously steep slopes), but did not attempt to identify other hazards like slippery sand since doing so would have been prohibitively difficult (Maimone, Yang and Matthies 2007).

MoonRanger has an existing definition of what obstacles should be identified, created for benchmarking camera specs. With the current wheel design, objects up to 8cm high can be traversed (MoonRanger 2019). As such, the decision was made that obstacles up to 5cm x 5cm x 5cm should be identified (Teza and Kumar 2019). Additionally, these obstacles should be able to be identified when they are between 1 and 2 meters from the rover as this is the optimal range to inform planning around them (Kumar 2019). Camera resolution has been specified such that these obstacles at these ranges will take up at least 50 pixels in the image (Teza and Kumar 2019).

Criterion 1, Objective 1: The final recommendation algorithm should provide sufficient information to identify obstacles 5 cm x 5cm x 5 cm at between 1 and 2 meters from the rover. These obstacles will occupy at least 50 pixels in the images used for stereo vision.

There are many ways to achieve this objective: a full 3D rendering of the terrain in front of the rover would provide sufficient information to detect these obstacles, but so could a few or even a single point of range information on each obstacle, theoretically. So this objective does not restrict what algorithms it is reasonable to study, but it does provide an objective that the algorithms must complete to be optimal for the final recommendation.

Providing Proof of Safe Routes

If the final recommendation algorithm identified all obstacles 100% of the time, then any route with no obstacles detected on it would implicitly be safe. Since it is unlikely or impossible to achieve this standard of obstacle detection in the final recommendation algorithm, another objective for the output results is to distinguish between routes that are safe and routes where safety is unknown.

No official MoonRanger specifications exist with regard to this objective as far as I am aware of, so it is not a requirement of any final recommendation algorithm. However, it is a desirable characteristic.

Criterion 1, Objective 2: The final recommendation algorithm should, where possible, distinguish regions that are free from obstacles and hence "proved safe" from regions where the presence or absence of obstacles is unknown.

This objective does restrict what algorithms it is reasonable to study, since only some algorithms provide a guarantee of attempting to prove safe paths clear of obstacles. Dense algorithms are likely best at this task, since they attempt to provide depth information for the entire image and generally deal well with poorly-textured regions (and safe paths will likely generally consist of poorly-textured dust). Quasi-dense and sparse algorithms with some guarantees about forcing regularly-spaced depth measurements across the entire image could also achieve this objective. Since this objective is desired but not a total requirement, some algorithms may be studied that are not expected to achieve this objective, but algorithms that are expected to achieve this objective are preferred.

Integration with Visual Odometry

Aside from detecting obstacles for route-planning the other primary use of stereo vision in MoonRanger is performing visual odometry. The goal of visual odometry is to provide reliable data on the rover's position on the surface of the moon that is not prone to IMU-drift or inaccuracies caused by wheel-slip.

The algorithm for visual odometry has not yet been decided, but currently an approach similar to VINS-Mono and VINS-Fusion is being considered (Kumar 2019). This and many other visual odometry implementations rely on feature matching and stereo matching of the same type that this project aims to study (Qin, Li and Shen, VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator 2018) (Qin, Cao, et al. 2019) (Maimone, Yang and Matthies 2007) (Li, et al. 2016). As such, it is likely that the output results of the final recommendation algorithm will be used as input to the visual odometry algorithm, and so it is desirable that they be compatible with that purpose.

Criterion 1, Objective 3: The final recommendation algorithm should output results that integrate easily into MoonRanger's chosen visual odometry algorithms.

Integration with Mapping

MoonRanger has decided not to pursue a simultaneous localization and mapping (SLAM) approach to mapping and navigation, since doing so with current algorithms is too computationally expensive (MoonRanger 2019). Nevertheless, there is ongoing discussion of the possibility that some of the data acquired from stereo vision during a traverse should be saved to the map as breadcrumbs to be used for navigation during the return to the lander. If this is to be the case, then it would be desirable for the output data from the final recommendation algorithm to be compatible with this task.

Criterion 1, Objective 4: If the decision is made to save output results from stereo vision to a map to help the rover with navigation, then the final recommendation algorithm should output results that are compatible with this task.

Criterion 2: Accuracy

Type of Accuracy

There are different ways to measure accuracy, and a simple approach is unlikely to be ideal for the purposes of MoonRanger. Recall that the main goal of stereo vision for obstacle avoidance is to detect 5 cm x 5cm x 5cm objects within 1 to 2 meters in front of the rover (Teza and Kumar 2019) (Kumar 2019). A simple measure of accuracy would be to compare the true depth of all pixels in the image to the depth computed by stereo vision. But this measure could heavily penalize a relatively small depth-inaccuracy over a large section of flat terrain (not much of a problem for the purposes of obstacle avoidance) and may not heavily penalize detecting a false obstacle or failing to detect a real obstacle since an obstacle could take up only approximately 50 pixels in the image (yet this would be a large problem for the purposes of obstacle avoidance). Useful accuracy in the context of MoonRanger should chiefly measure whether obstacles were correctly detected or not; other aspects of how well the result correlates with the ground truth are less important.

Additionally, false negatives (failure to detect an obstacle when one actually exists) are generally worse than false positives (detecting an obstacle when one does not actually exist).

Criterion 2, Objective 1: Accuracy should be assessed in a way consistent with the purpose of the stereo vision algorithm: detecting obstacles. The final recommendation algorithm should be considered accurate if it primarily has a low rate of false negatives and secondarily a low rate of false positives in detecting obstacles. Other aspects of accuracy should be tertiary.

Operational Conditions Affecting Accuracy

The operational conditions of the lunar polar environment pose unique challenges for stereo vision. Finding the best solution to these challenges is the purpose of this assignment. As such, the assessment of accuracy for the final recommendation should reflect these operational conditions.

The following is a list of operational conditions that should be incorporated in assessing accuracy. This list is likely incomplete and should be expanded as the project continues:

- The effects of the lunar environment on imagery:
    - The absence of atmospheric effects on lighting and imagery
    - The effects of how the lunar regolith scatters light (Wong, et al. 2017)
    - The “appearance” of lunar terrain (monochrome, large regions of dust, etc.)
- The effects of the polar lunar environment on imagery:
    - Low sunlight angle and long shadows
- The camera types and configuration to be used by MoonRanger
- Whether objects are in the 1-2 meter detection range preferred by MoonRanger
- Whether objects are detected constantly, sometimes, or never while they pass through the rover’s field of view

No dataset can perfectly capture all of these conditions simultaneously, since there are no surface images from the moon’s polar regions (let alone surface images from the moon’s polar regions using the same cameras as MoonRanger in the same configuration as MoonRanger). So a combined approach is likely to be the best possible.

NASA’s POLAR Stereo Database is the best resource for images that include the operational conditions of lunar polar regions, and it is likely that the images can be processed to put them in a form similar to that which MoonRanger’s cameras will deliver (resolution, color space, etc.). Still, there are operational conditions of the MoonRanger rover that this database cannot capture. For example, if an algorithm frequently failed to detect obstacles at a range of 2 meters, but would reliably detect them as the rover approached to within 1.5 meters, this would not be ideal but it would be a pretty good result nonetheless. The POLAR database cannot check for this behavior, though: it provides only a small number of static images, and so one cannot test whether the obstacles in the images would have been detected as the camera moved closer. Detection of obstacles in the POLAR images is all or nothing. Conversely, data from test drives with CMU autokrawlers can mimic operational conditions like the camera specifications and driving (camera motion) behavior expected from MoonRanger well, but this data does not mimic the lunar and lunar polar conditions nearly as well.

Criterion 2, Objective 2: Tests of accuracy should include as many of the operational conditions expected during the actual MoonRanger mission as possible. Multiple test datasets should be used to try to assess *how all operational conditions effect accuracy if it is not possible to mimic all operational conditions in one dataset.*

Depth vs Breadth of Study

One final objective to consider related to accuracy is depth versus breadth in the choice of algorithms to study. There is a finite time available for this project, and studying one algorithm can take longer or less long depending on how difficult to implement the algorithm is and how much time is spent tweaking and refining the parameters and techniques used in the algorithm to maximize performance. Breadth, in the form of picking algorithms that are quick to assess, is desirable since without it the best algorithms may be overlooked. But depth is also important since the algorithms with the best performance may be complicated and/or recently-made (and hence lacking an open-source implementation, greatly increasing the time needed to implement the algorithm before assessing it), and may need tweaking and refinement of parameters to get the best performance. This project should work to find a balance between these competing demands.

Criterion 2, Objective 3: The set of algorithms evaluated by this study should balance the desire for a broad survey of many algorithms with the desire for a deep investigation of algorithms that may be slow to implement. The goal of this balancing is for the final recommendation algorithm found to have the best performance possible in all criteria and objectives.

Criterion 3: Efficiency

Notes in cyan highlights are last-minute revisions based on discussion during ANS team meeting on 12/14 planning for spring semester.

Efficiency is highly desirable. The question of what computing hardware will be available on MoonRanger is not currently settled (although it has been identified as a high priority to resolve this question as soon as possible), but whatever hardware is available will be limited. It is possible that a design with a GPU (Nvidia TX1, TX2, or TK1) will be selected, but it is also possible that the only computer onboard will be a 43 Mhz ARM Cortex-M3 processor (MoonRanger 2019). It currently seems most likely that if a single computer layout were selected, the Xilinx Zynq-7020 System on Chip would be chosen, which runs at 766Mhz (J. Teza 2019) (Space Micro 2019). Currently a frame rate of 1 FPS is targeted and the goal is for all of the navigation software to run in real-time enabling continuous motion. Since other expensive algorithms like path planning and visual odometry will also be running, the stereo vision algorithm should be able to run to completion on an image pair in a few hundred milliseconds at most. Running even faster than this is desirable, as the spare time created by doing so could be used as a safety margin, or given to other algorithms that need more time, or could enable the computer to be run in a lower-energy state.

Criterion 3, Objective 1: The final recommendation algorithm should be as efficient as possible.

This objective is similar to criterion 1 objective 2 in that it also somewhat restricts what algorithms are reasonable to study. Interestingly, it does so in almost the exact opposite way. Criterion 1 objective 2 encouraged dense or at least quasi-dense algorithms. But these algorithms are usually the most computationally expensive, whereas sparse algorithms can be much more efficient. The balancing act between criterion 1 objective 2 and this objective will likely be one of the biggest challenges in this study.

Project Goals

Overview

Notes in cyan highlights are last-minute revisions based on discussion during ANS team meeting on 12/14 planning for spring semester.

The previous section described the criteria by which the usefulness of the final recommendation algorithm could be measured. These are not the criteria by which the successfulness of this project should be judged. Rather, the project's goal is to deliver the final recommendation algorithm best meeting the above criteria. This project will be a success if it does indeed deliver a final recommendation algorithm close to ideal (or some other equivalently useful results), and a failure if it delivers a recommendation that is far from the best possible.

The goal of this project is to contribute an algorithm that is useful to the MoonRanger project, which does introduce other logistical challenges, though. ~~There is no clear current desired timeline for the implementation of the stereo vision algorithm (and this project should ideally conclude just before implementation of the stereo vision algorithm starts).~~ This project will implement the stereo vision algorithms studied in such a way that they can work inside the MoonRanger pipeline. At any point in time, the currently-best algorithm of those studied (that is implemented with sufficient stability) will be the algorithm used in the MoonRanger pipeline. Additionally, there are no clear current desired timelines for the development of the visual odometry algorithm or the mapping and navigation algorithms. There are good reasons for both of these algorithms to be closely tied to the stereo vision algorithm, and so their development and this study should probably be integrated and inform each other, which further influences when it is ideal for this study to occur.

As such, the original goals and milestones for this project from the October proposal are included here (slightly revised based on new information) as an example of what the goals and milestones for the project should look like assuming that the study should be completed at the end of Spring 2020. If some earlier completion date for the study becomes preferable during Spring 2020 semester, this list of goals and milestones will be revised to complete the study faster and add additional goals to do with implementation of the stereo vision algorithm.

Current Goal List

At the 100% level of completion, the final goals for this project are as follows. The "edge 2" algorithms/pipelines for stereo vision will be evaluated to gauge their accuracy in lunar imagery, and their efficiency in a running environment representative of that on MoonRanger. The "edge 2" algorithms are defined as follows: the algorithm that is expected to have the highest level of accuracy, and the algorithm that is expected to be the most efficient. In addition to these "edge 2," 5 more algorithms/pipelines shall be tested in the same way. These will be chosen from the research to maximize expected quality of the final solution found and the breadth of the solutions considered. After these tests are complete, 2 algorithms/pipelines will be created, possibly as hybrids of those already tested, with the goal being for one of them to be the best solution overall. These 2 final algorithms will be evaluated and tweaked to maximize performance, and a final recommendation will be presented, as well as the results for all algorithms.

These goals imply some necessary stepping-stone goals. Namely, conducting a review of the existing research to identify promising algorithms, and designing and building testing setups for accuracy and efficiency that mirror the environment in which MoonRanger will operate.

If this project proceeds faster than expected, more algorithms will be evaluated. The 125% level of success will see evaluations of 7 algorithms/pipelines from the research in addition to the "edge 2" (as opposed to 5 at the 100% level), as well as 3 final hybrid algorithms (as opposed to 2).

If this project encounters challenges or proceeds slower than expected, there are two alternate 75% level of success goal sets. The two choices let this project adjust to avoid the largest challenges and/or maximizes focus on the most useful parts of the project.

The first 75% level of success is a simple reduction in the number of algorithms evaluated. Only the "edge 2" and 3 other algorithms will be evaluated, and only 1 final hybrid algorithm will be created (as opposed to 5 and 2 respectively at the 100% level).

The other 75% level of success will see the "edge 2" and 7 other algorithms evaluated as well as 3 final hybrid algorithms created (this is the same as in the 125% success level). However, efficiency evaluation will not be performed on any algorithm. At this level of success, the project will only assess the accuracy of the algorithms.

Current List of Algorithms to Study

This list shows the current algorithms considered most promising for study in this paper. This list will almost certainly be updated during the spring semester as results emerge from the first testing and inform the best way to proceed and as the goals and plans of other aspects of MoonRanger change.

Notes in cyan highlights are last-minute revisions based on discussion during ANS team meeting on 12/14 planning for spring semester.

"Edge 2" Algorithms

- Expected-Most-Accurate: ASIFT-based stereo pipeline similar to that in *Stereo Vision Technologies for China's Lunar Rover Exploration Mission* (2016) by Mingli Li, Shaochuang Liu, Youqing Ma, Hao Ma, Sun Zezhou, Jia Yang, and Changming Sun and *A Stereo Matching Algorithm for Lunar Rover* (2010) by Fengping Cao and Rongben Wang
- Expected-Most-Efficient: ExFast-based stereo as described in *A New Feature Detector and Stereo Matching Method for Accurate High-Performance Sparse Stereo Matching* (2012) by Konstantin Schauwecker, Reinhard Klette, and Andreas Zell

Other Algorithms to Study:

1. The stereo technology used by VINS-Mono as described in *VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator* (2017) by Tong Qin, Peiliang Li, and Shaojie Shen
2. High-performance and tunable stereo reconstruction algorithm as described in *High-Performance and Tunable Stereo Reconstruction* (2016) by Sudeep Pillai, Srikumar Ramalingam, and John J. Leonard
3. Two-step expansion dense stereo as described in *Feature Based Stereo Matching Using Two-Step Expansion* (2014) by Liqiang Wang, Zhen Liu, and Zhonghua Zhang
4. The Mars Exploration Rover Stereo Vision algorithm as described in *Two years of Visual Odometry on the Mars Exploration Rovers* (2007) by Mark Maimone, Yang Cheng, and Larry Matthies

5. Optimized implementation ELAS as described in *Real-Time Dense Stereo Matching With ELAS on FPGA-Accelerated Embedded Devices* (2018) by Oscar Rahnama, Duncan Frost, Ondrej Miksik, and Philip H.S. Torr and *Efficient Large-Scale Stereo Matching* (2011) by Andreas Geiger, Martin Roser, and Raquel Urtasun or Lib-ELAS
6. The stereo vision component of ORB-SLAM as described in *ORB-SLAM: A Versatile and Accurate Monocular SLAM System* (2015) by Raul Mur-Artal, J. M. M. Montiel, and Juan D. Tardos
7. OpenCV basic stereo vision library (2 different algorithms)

Milestones (at the 100% level of success)

The bi-weekly milestones are as follows. Note that any time testing an algorithm is mentioned, implementing this algorithm is implicitly assumed to be part of the milestone goal. “Implementing” in this context may be as simple as downloading and running an open-source implementation.

Notes in cyan highlights are last-minute revisions based on discussion during ANS team meeting on 12/14 planning for spring semester.

27th January: ~~build accuracy testing system based on POLAR database and evaluate expected-most-efficient algorithm on it~~ implement expected-most-efficient algorithm and integrate it with the overall MoonRanger pipeline evaluate its accuracy on test autokrawler test data if such data exists (might be ad-hoc accuracy evaluation to begin with; permanent automated accuracy evaluation system to be decided later)
10th February: ~~build accuracy testing system based on autokrawler data and evaluate “edge 2” algorithms on both accuracy tests~~ build accuracy testing system based on POLAR database and evaluate “edge 2” algorithms on both accuracy tests
24th February: build efficiency testing system, evaluate efficiency of “edge 2” algorithms
16th March: evaluate efficiency and accuracy of 2 additional algorithms
30th Match: evaluate efficiency and accuracy of 3 further additional algorithms
13th April: implement and fine-tune the 2 final hybrid solutions
27th April: evaluate efficiency and accuracy of the 2 final hybrid solutions, make final recommendation about best overall algorithm, present results and discussion on testing of all algorithms

Literature Search

In the October proposal, I identified four[1] goals of the literature search. In this version of the proposal I report my progress and future intentions with regard to each of the four goals.

Context for Autonomous Planetary Rovers

This goal is to read and study literature relating to autonomy for planetary rovers in order to understand the context of the MoonRanger project. This context should then inform my selection of project topic and how I conduct my project. This goal was identified as “mostly complete” in October, and remains “mostly complete.” I have learned sufficient context about the project to develop this project proposal, and will continue to learn more context as I work with the MoonRanger team further. I have read the following papers in researching for this goal:

[1] In the October proposal, I split the “Best Methods for Evaluating Accuracy and Efficiency” goal into two goals and so actually had five goals. In this proposal, I find that spilt unnecessary and treat the two goals as one.

- *Developing a Framework for Reliable Autonomous Surface Mobility* (2012) by David Wettergreen and Michael Wagner
- *Science Autonomy for Rover Subsurface Exploration of the Atacama Desert* (2014) by David Wettergreen, Greydon Foil, Michael Furlong, and David R. Thompson
- *Strategic Autonomy for Reducing Risk of Sun-Synchronous Lunar Polar Exploration* (2018) by Nathan Otten, David Wettergreen and William Whittaker
- *D* Lite* (2002) by Sven Koenig and Maxim Likhachev
- *Field Experiments in Robotic Subsurface Science with Long Duration Autonomy* (2017) by Srinivasan Vijayarangan, David Kohanbash, Greydon Foil, Kris Zacny, Nathalie Cabrol and David Wettergreen
- *A Stereo Matching Algorithm for Lunar Rover* (2010) by Fengping Cao and Rongben Wang
- *Stereo Vision Technologies for China's Lunar Rover Exploration Mission* (2016) by Mingli Li, Shaochuang Liu, Youqing Ma, Hao Ma, Sun Zezhou, Jia Yang, and Changming Sun
- *Two years of Visual Odometry on the Mars Exploration Rovers* (2007) by Mark Maimone, Yang Cheng, and Larry Matthies

I also met with David Wettergreen on the 24th of October to learn about the MoonRanger project, and have attended several MoonRanger team meetings including the MoonRanger Ideation Review on November 19th.

Identification of the Best Algorithms

This goal is to study the existing research on stereo vision algorithms in order to find the best algorithms and techniques to include in this study. In October, this goal was "not yet started." At present, it is "mostly complete." I have read papers on a large number of diverse approaches to stereo vision and I believe this research forms a solid basis from which to begin this study. As the study progresses, new discoveries by me of which approaches work and which don't and new decisions by MoonRanger of what directions and algorithms to peruse will likely lead to further research on this topic.

- *A Stereo Matching Algorithm for Lunar Rover* (2010) by Fengping Cao and Rongben Wang (also listed in previous category)
- *Stereo Vision Technologies for China's Lunar Rover Exploration Mission* (2016) by Mingli Li, Shaochuang Liu, Youqing Ma, Hao Ma, Sun Zezhou, Jia Yang, and Changming Sun (also listed in previous category)
- *Two years of Visual Odometry on the Mars Exploration Rovers* (2007) by Mark Maimone, Yang Cheng, and Larry Matthies (also listed in previous category)
- *Review of visual odometry: types, approaches, challenges, and applications* (2016) by Mohammad O. A. Aqel, Mohammad H. Marhaban, M. Iqbal Saripan, and Napsiah Bt. Ismail
- *Efficient Large-Scale Stereo Matching* (2011) by Andreas Geiger, Martin Roser, and Raquel Urtasun
- *High-Performance and Tunable Stereo Reconstruction* (2016) by Sudeep Pillai, Srikumar Ramalingam, and John J. Leonard
- *Real-Time Dense Stereo Matching With ELAS on FPGA-Accelerated Embedded Devices* (2018) by Oscar Rahnama, Duncan Frost, Ondrej Miksik, and Philip H.S. Torr
- *Feature Based Stereo Matching Using Two-Step Expansion* (2014) by Liqiang Wang, Zhen Liu, and Zhonghua Zhang
- *A New Feature Detector and Stereo Matching Method for Accurate High-Performance Sparse Stereo Matching* (2012) by Konstantin Schauwecker, Reinhard Klette, and Andreas Zell

- *VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator* (2017) by Tong Qin, Peiliang Li, and Shaojie Shen
- *A General Optimization-based Framework for Global Pose Estimation with Multiple Sensors* (2019) by Tong Qin, Shaozu Cao, Jie Pan, and Shaojie Shen

There are additional papers in this category that I intend to read in the near future, but have not yet read:

- *ORB-SLAM: A Versatile and Accurate Monocular SLAM System* (2015) by Raul Mur-Artal, J. M. M. Montiel, and Juan D. Tardos
- *Direct Sparse Odometry* (2017) by Jakob Engel, Vladlen Koltun, and Daniel Cremers
- *Visual-inertial Monocular SLAM with Map Reuse* (2017) by Raul Mur-Artal, and Juan D. Tardos

Best Methods for Evaluating Accuracy and Efficiency

This goal is to discover the best practices to use for assessing the accuracy and efficiency of different stereo vision algorithms. In October, this goal was "not yet started." At present, it is "mostly complete." In reading through numerous papers on stereo vision, I did not find any common detailed set of guidelines and procedures for evaluating algorithms. In general, the algorithms were tested on a database of relevant images and if efficiency evaluation was desired were tested on streaming video data while running on specific hardware. I plan to follow this general setup in my study (although I may substitute virtual emulation or use of similar hardware instead of use of exact hardware if I find doing so does not affect the accuracy of my efficiency measurements too greatly).

Finding Lunar Imagery Data to Use for Evaluation

This goal is to acquire the data to use for testing the accuracy of the algorithms I study (read section ***Operational Conditions Affecting Accuracy*** to see what factors inform what data should be used). In October this goal was "partially complete." In October, I had examined the data available online from the Apollo, Lunokhod, and Yutu 1 missions. I was also aware of the NASA POLAR database, but had not yet examined it. Since then, I have examined the NASA POLAR database and have also become aware of test data collected by MoonRanger driving Autokrawlers in environments similar to the lunar surface. I still have not investigated whether there is data available online from the Yutu 2 mission. I believe I have adequate data to use for this study at this point, and so this goal is "mostly complete." Remaining work on this goal includes searching for Yutu 2 data, examining MoonRanger Autokrawler data, and possibly aiding in identifying new scenarios for which Autokrawler data should be gathered.

Resources Needed

Research Papers

Access to some research papers will be obtained through CMU libraries access to online research databases. This access is freely available to CMU students.

No additional action is needed in this category to acquire resources.

Algorithms and Software

MoonRanger software team currently works in ROS, which is open source and freely available online. To run Linux code in a virtual machine I will use either freely available software (Windows Subsystem Linux) or software provided free to CMU students (VMWare). All algorithms I test in this study will either be from free open source software and/or written by me based on descriptions in research papers.

No additional action is needed in this category to acquire resources.

Testing Databases

NASA's POLAR database is freely accessible online. Data from MoonRanger Autokrawlers is available to all members of MoonRanger. New data with specific scenarios or configurations of MoonRanger Autokrawlers may be acquired during this project if it is useful to this project and MoonRanger at larger. It will not compromise this project if such data is not acquired.

MoonRanger hardware may be used during Spring semester (gathering additional test data from the Autokrawlers) to improve this project, but such use is not essential for this project.

Hardware for Efficiency Testing

Notes in cyan highlights are last-minute revisions based on discussion during ANS team meeting on 12/14 planning for spring semester.

It is currently unknown whether testing the efficiency of algorithms can be done effectively on my personal laptop or if running on the same computers as planned for the mission will be needed. Either way, some tests on the computers planned for the mission will be needed if only to demonstrate that laptop-based efficiency tests are accurate. Such testing is impossible until the choice of computers planned for the mission is decided.

MoonRanger hardware may be used during Spring semester (testing efficiency on computers of the type planned for the mission). This usage is important for the project, although most of the project goals could be achieved without it.

MoonRanger team is expecting to test the entire mapping and navigation pipeline on surrogate copies of the mission computers during spring semester; this testing will include testing the code I write.

References

Kumar, Varsha, interview by William Lee-Moore. 2019. (December 5).

Li, Mingli, Shaochuang Liu, Youqing Ma, Hao Ma, Sun Zezhou, Jia Yang, and Changming Sun. 2016. "Stereo Vision Technologies for China's Lunar Rover Exploration Mission." *International Journal of Roboitcs and Automation* 128-136.

Maimone, Mark, Chen Yang, and Larry Matthies. 2007. "Two years of Visual Odometry on the Mars Exploration Rovers." *Journal of Field Robotics* 169-186.

MoonRanger. 2019. *MoonRanger Ideation Review.* November 19.

Qin, Tong, Peiliang Li, and Shaojie Shen. 2018. "VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator." *IEEE Transactions on Robotics* 1004-1020.

Qin, Tong, Shaozu Cao, Jie Pan, and Shaojie Shen. 2019. "A General Optimization-based Framework for Global Pose Estimation with Multiple Sensors." *arXiv.*

Space Micro. 2019. *Cubesat Space Processor.* Product Datasheet, San Diego: Space Micro Sales.

Spice, Byron. 2019. "NASA Selects Carnegie Mellon, Astrobotic To Build Lunar Robot." *CMU News*, July 3.

Teza, James. 2019. "MoonRanger Avionics Meeting 20191204." December 04.

Teza, Jim, and Varhsa Kumar. 2019. *MoonRanger Software Slack.* November 12.

Wettergreen, David, and Michael Wagner. 2012. "Developing a Framework for Reliable Autonomous Surface Mobility." *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space.*

Wettergreen, David, Greydon Foil, Michael Furlong, and David Thompson. 2014. "Science Autonomy for Rover Subsurface Exploration of the Atacama Desert." *AI Magazine* 35 (4): 47-60. http://search.proquest.com/docview/1644634582/.

Wong, Uland, Ara Nefian, Larry Edwards, Xavier P Buoyssounouse, Michael Furlong, Matt Deans, and Terry Fong. 2017. "Polar Optical Lunar Analog Reconstruction (POLAR) Stereo Dataset." NASA Ames Research Center, May.