Efficient Stereo Vision for Feature-sparse Lunar Images

15-300 Project Proposal

William Lee-Moore

PROJECT WEBSITE

The website for this project is at: www.andrew.cmu.edu/user/wleemoor/lunar-vis-project.html

PROJECT DESCRIPTION

My project will be a collaboration with the ongoing CMU project MoonRanger. MoonRanger is a project intending to put an autonomous robotic rover, called MoonRanger, on the moon (Spice 2019). This 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. I will attend the regular team meetings for the teams relevant to my project (predominantly the mapping sub-team of the autonomy team) as my primary method of contacting and coordinating with the efforts of the project, as well as maintaining communication with a faculty mentor to keep track of project progress and receive advice/guidance as necessary. It currently appears most likely that Dr. Wettergreen will be my faculty advisor, but this has not been fully finalized yet.

In general, this project proposal represents a preliminary idea for the project direction. This proposal is based on a project idea suggested by Dr. Wettergreen during a meeting on 10/24. However, the needs of the MoonRanger project are still changing quickly, and based on further conversation with students in the team, it appears that the project may want to commit to a vision algorithm and pipeline *during* spring 2020 semester, which would render a recommendation on the ideal vision algorithm and pipeline *at the end of* spring 2020 semester redundant. If this is indeed the case, the project will be transformed to align better with what would create a useful contribution for MoonRanger. For example, the project may feature a briefer phase of selecting the vision algorithm and pipeline and have a longer phase of tweaking and refining its performance added. As such, the discussion of MoonRanger and the challenges of stereo vision in lunar images is all fully accurate, but the discussion of the project's goals and milestones is only accurate as an example of the general sort of expected project purpose, work, and goals. The details of what exact deliverables are desired and how the milestones are arranged to achieve them may change greatly.

This project will enhance the accuracy and speed of the stereo vision used by MoonRanger, thereby making the rover capable of safer and faster operations. To map the terrain in front of it for navigation, MoonRanger uses stereo vision: matching features between images taken at different locations to triangulate those features in 3D space and hence generate a point cloud of the scene in the images. Feature matching can be difficult in lunar images, primarily because of large regions of featureless dust, but also because of the lack of color reducing distinctiveness of all features and long shadows hiding large regions of the image. There exist various techniques, for example Random Sample Consensus (RANSAC) (Narasimhan 2019), that can detect and remove false matches, but the mere absence of correct matches can be a large problem. If the rover cannot match features in a large region of dust ahead of it, it will not have a point cloud to map this region, and must choose to either withdraw or drive into the unknown. Neither choice is desirable. As such, my project will analyze existing research into these stereo vision challenges to attempt to find the best algorithms to generate as many accurate feature matchings as possible given the challenges of lunar imagery and the other challenges intrinsic to MoonRanger.

MoonRanger presents other intrinsic challenges. One already discussed was the changing nature of the MoonRanger project will require this project to adapt its goals and timeline dynamically and/or take other measures to work around uncertain specifications and deliver useful results despite fuzzy requirements. For example, the number of cameras and whether a graphics card will be on the rover are both still undecided. Another bigger challenge is that MoonRanger will have limited computation power, so all algorithms used must be highly efficient. There is often a natural tradeoff between efficiency and accuracy in computer vision tasks. For example, the BRIEF feature descriptor can be up to 41 times faster to compute than SURF, but BRIEF fails to match features that have been significantly rotated in the image whereas SURF does not (Calonder, et al. 2010). Additionally, the large volume of solutions in existing research and the fact that the optimal solution for MoonRanger may require a hybrid of existing algorithms also make it more challenging to find the ideal solution.

The primary contribution of this project is to improve the expected success achieved by the MoonRanger project. MoonRanger's main contribution will be demonstrating how autonomy can make space rovers capable of executing advanced missions away from their landers without needing heavy and expensive radios to maintain constant communication with controllers on Earth. Additionally, however, other projects facing a need for efficient stereo vision in feature-sparse situations or other related challenges may benefit from the comparison of algorithm accuracy and efficiency produced by this project.

PROJECT GOALS

At the 100% level of completion, the final goals for this project are as follows. The "core 3" 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 "core 3" algorithms are defined as follows: the current algorithm MoonRanger uses, 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 "core 3," 4 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 6 algorithms/pipelines from the research in addition to the "core 3" (as opposed to 4 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 "core 3" and 2 other algorithms will be evaluated, and only 1 final hybrid algorithm will be created (as opposed to 4 and 2 respectively at the 100% level).

The other 75% level of success will see the "core3" and 6 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.

MILESTONES (AT THE 100% LEVEL OF SUCCESS)

For the first technical milestone for 15-300, to occur at the end of the Fall 2019 semester, the following shall have been accomplished. A review of the existing research and literature into methods for stereo vision relevant to these challenges will have been conducted, and from this research a list of 8 algorithms to evaluate will have been compiled: the expected most accurate, the expected most efficient, 2 more to evaluate at 75% success, 2 in addition to that for 100% success, and 2 in addition to those for 125% success. Additionally, the plan for testing setups for accuracy and efficiency will have been developed, guided by the review of related research, meetings with the relevant MoonRanger teams, meetings with Dr. Wettergreen, and meetings with other faculty or students at CMU who have useful experience on how to construct such tests. Also, all necessary resources for the testing setups will have been assembled, or at least ordered (see resources needed section).

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.

27th January: build accuracy testing system and calculate any necessary ground-truth results for it
10th February: evaluate accuracy of current and expected-most-accurate algorithms
24th February: build efficiency testing system, evaluate accuracy and efficiency of expected-fastest algorithm, evaluate efficiency of current and expected-most-accurate algorithms
16th March: evaluate efficiency and accuracy of 2 additional algorithms
30th Match: evaluate efficiency and accuracy of 2 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

There are five essential questions to be answered by the literature search of which one is mostly complete at this time, one is partially complete, and three are still to be started. The mostly complete task is reviewing literature regarding autonomy of planetary rovers for context about the challenges faced by MoonRanger. The still-to-be-started tasks are identification of the best algorithms and approaches for stereo vision under these circumstances (I have some familiarity with this topic already, but have not yet started a dedicated search for this project), and identification of the best way to evaluate the accuracy and efficiency of stereo vision algorithms. The partially completed task is finding the lunar imagery data to use for evaluation: I have found workable images from the Apollo program, excellent images from the Chinese Yutu rover (Chinese Academy of Sciences / China National Space Administration / The Science and Application Center for Moon and Deepspace Exploration (ed. Lakdawalla, Emily) 2016), and also achieved reasonably high confidence that there are not high quality images available from the Soviet Union's Lunokhod rovers. I do not yet know if there exist good images available from the Yutu 2 rover, or good simulated images from NASA or other space agencies.

RESOURCES NEEDED

For this project, I will need to implement all algorithms I need to test in some way, meaning an existing implementation or detailed description will be needed. However, it is likely that all algorithms I wish to evaluate will either have an open-source implementation already developed or a clear enough description for me to implement the algorithm myself. For algorithms where this is not the case, it is likely an alternate algorithm can be selected that is nearly as good for the purposes of this project.

A database of lunar images for the accuracy evaluation is also needed. I have already discovered that sufficient databases exist open-source online.

Finally, a method of testing efficiency in a way that generates results applicable to MoonRanger is needed. Hopefully there is a way to do this through emulation or through running on a standard with restricted settings. If this does not generate results that are accurate/applicable for the MoonRanger hardware, though, running time on that (or equivalent) hardware may be needed for the efficiency tests.

REFERENCES

Calonder, Michael, Vincent Lepetit, Christoph Strecha, and Pascal Fua. 2010. "BRIEF: Binary Robust Independent Elementary Features." *European Conference on Computer Vision.* Crete: Springer. 778-792.

Chinese Academy of Sciences / China National Space Administration / The Science and Application Center for Moon and Deepspace Exploration (ed. Lakdawalla, Emily). 2016. "Chang'e 3 data: Rover Panoramic Camera (PCAM)." *The Planetary Society.* Jaunary 28. Accessed October 26, 2019. https://planetary.s3.amazonaws.com/data/change3/pcam.html.

Narasimhan, Srinivasa. 2019. "Image homographies." *16-385: Computer Vision.* October 16. Accessed October 22, 2019.

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

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/.