Appears in Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation for Space,
Brisbane, Australia, November 2024

Onboard Autonomous Health Assessment and Global Localization for the Mars Helicopter: Towards Multi-Flight Operations

[1]C. Basich, [1]C. Mauceri, [1]G. Kubiak, [1]J. Delfa, [1]A. Candela-Garza, [1]Pedro Proença, [1]Barry Ridge, and [1]S. Chien

[1]Jet Propulsion Laboratory, California Institute of Technology, USA

ABSTRACT

NASA's Mars Helicopter, *Ingenuity*, performed the first successful powered, controlled flight on another planet on April 19, 2021, marking a historic moment in the advancement of space robotics. Since then, Ingenuity performed 72 successful flights over its lifespan, well over 10 times what it was originally planned to do. Part of Ingenuity's longevity is due to an extensive post-flight, ground-in-the-loop cycle in which the system's health, flight-readiness, and location are rigorously ascertained from expert analysis of downlinked data by human engineers. However, this process not only requires the time and attention of engineers, it also prevents the helicopter from engaging in multi-flight operations, sometimes for days, severely limiting its potential science productivity. In this paper, we present our recent effort developing a fully onboard, automated health assessment and global localization software for the spacecraft. We discuss the challenges presented by the problem setting itself, the mission parameters that influenced our approach and evaluations, and present initial results from both contributions in the context of Ingenuity, which flew the onboard health assessment software in January 2024.

INTRODUCTION

The potential for the aerial mobility of rotorcraft on Mars demonstrated by Ingenuity fundamentally alters and expands the space of science missions that can be investigated in the future. Aerial platforms can cover multiple dimensions of exploration (as opposed to rovers which are constrained to flat, even terrain), expanding the range potential of exploration systems, enabling access to scientifically interesting but hazardous or otherwise unreachable terrains, and providing a unique opportunity for in-situ atmospheric science at different elevations. These capabilities are opportunities to address high-priority science questions identified in the Mars science community, demonstrating the vast potential of these systems in the future.

To ensure the health and success of the helicopter over the course of its deployment, Ingenuity's operational procedure included a ground assessment cycle after each flight to manually validate the system's health and flight-readiness, and globally localize the helicopter on Mars. Although the ground assessment cycle was effective at ensuring Ingenuity's health, the necessity of performing a ground-in-the-loop cycle post-flight necessarily [1, 2] precluded the possibility of multi-flight operation, i.e. the autonomous execution of multiple flights by the system without ground involvement, limiting the potential scientific productivity of the helicopter. Autonomous multi-flight operation would enable greater science productivity in future missions not only by increasing the operational efficiency of the system, but also by enabling more intricate maneuvers and flight plans, such as multiple hops to reach challenging but scientifically interesting terrains and locations, as well as spontaneous decision-making in the presence of unexpected but scientifically interesting phenomena.

Consequently, in order to enable future spacecraft to engage in autonomous multi-flight operation, we have developed an autonomy framework aiming to one day replace the full ground assessment cycle during nominal operation. Our primary contributions are therefore (1) our automated health assessment framework, OUTLAST (Onboard aUTomated heaLth ASsessmenT), which flew onboard Ingenuity in January 2024, and (2) our Automated Global Localization system, AGL. We present the results from initial testing of both frameworks in the context of Ingenuity, before discussing our work in the context of future space missions.

Onboard Automated Health Assessment (OUTLAST)

Description

OUTLAST (depicted in Figure 1) runs immediately post-flight for two key reasons. First, it ensures that the execution of the software will not impact the helicopter's performance during flight in any way. Second, compared to traditional, reactive FDIR methods [3, 4], operating post-flight enables the health assessment software to make decisions that are based on the entirety of the data collected during flight rather than just parts of it, enabling more comprehensive health assessments.

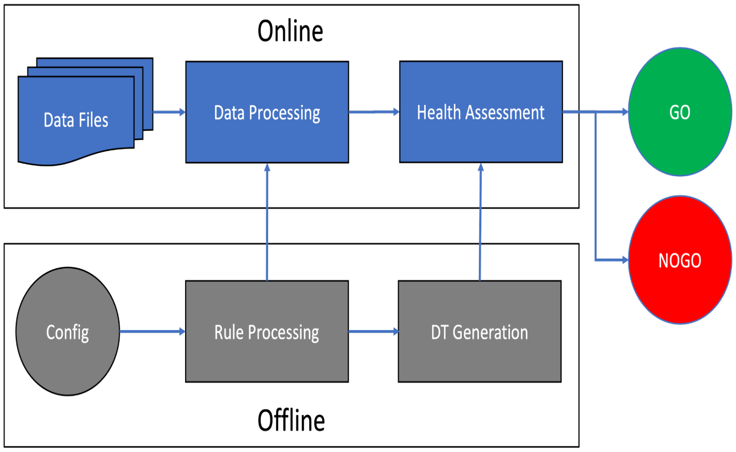

Figure 1: An illustration of the OUTLAST architecture.

Before OUTLAST is run onboard, the decision tree itself is produced on the ground; this decision tree may be produced by learning from data and/or from human-specified rules. Onboard, OUTLAST first reads into memory and processes all telemetry and flight logs from the preceding flight, and then prunes the data for any broken entries. After processing the data, a health assessment decision model is run onboard on the full dataset, and the results are compiled to (1) an GO/NOGO file and (2) a diagnostic file that provides an additional, compact source of information that can be downlinked immediately to be reviewed by the ground engineers. In our case, we used a decision tree [5] for our decision model as it it satisfied several important operational parameters: sample efficiency, explainability and interpretability (a key requirement), and the ability to encode both manually-specified decision criteria, as well as data-derived decision criteria (i.e., from training on prior flight data). However, other models could likewise be used, including random forests, gradient-boosted trees, gaussian processes, etc., and future work aims to explore these.

Results

We tested OUTLAST in several stages; first, we compiled the data from a total of seven prior flights to which we had access to complete datasets. Three of these flights were known to have been anomalous flights (i.e., flights where we expect a NOGO evaluation), and the other four were considered nominal flights, (i.e., flights where we expect a GO evaluation). OUTLAST correctly classified all of the off-nominal flights, producing valuable diagnostic data in the process that summarized all data channels on which the flight failed the health assessment. For the four nominal flights, one passed all decision criteria with no flags; two passed with one data channel failing, but it was determined to be within acceptable margins for preliminary testing to validate the approach ($0.082 < 0.083$, $0.088 > 0.085$ respectively), and one failing correctly despite its prior nominal status. In the last case, Ingenuity was found to have flown with a horizontal speed of 1.47 m/s during its descent, almost three times the nominal limit of 0.5 m/s. This case was uniquely important as it identified a case where off-nominal behavior was detected by OUTLAST in what was previously considered to be a completely nominal flight. This kind of diagnostic information can quickly provide valuable insights to the ground operators to inform them of a situation where further investigation should be performed prior to the following flight.

OUTLAST was successfully flown on Ingenuity on January 05, 2024, just prior to the helicopter's end of life, further supporting the potential for this approach on future spacecraft.

Onboard Automated Global Localization (AGL)

Description

Upon waking, Ingenuity had no sense of its position or orientation. Consequently, Navigation was performed in a frame of reference defined by its pose at start, meaning that the specification of waypoints by the operators is critical. However, this process itself could be challenging, relying on the images taken by the navigation camera just prior to landing to globally localize the helicopter. The global localization process took a navcam (navigation camera) image (see Figure 2, right) and a reference map [6] as input. The map has a resolution of 0.25 m/pixel and an orthographic projection. For Ingenuity, we reprojected the map to have a projection center at the Mars 2020 landing site (77.45088572 lon, 18.44462715 lat) and used the helicopter's VIO [7] estimated pose and the transformations between the helicopter body and the camera center.

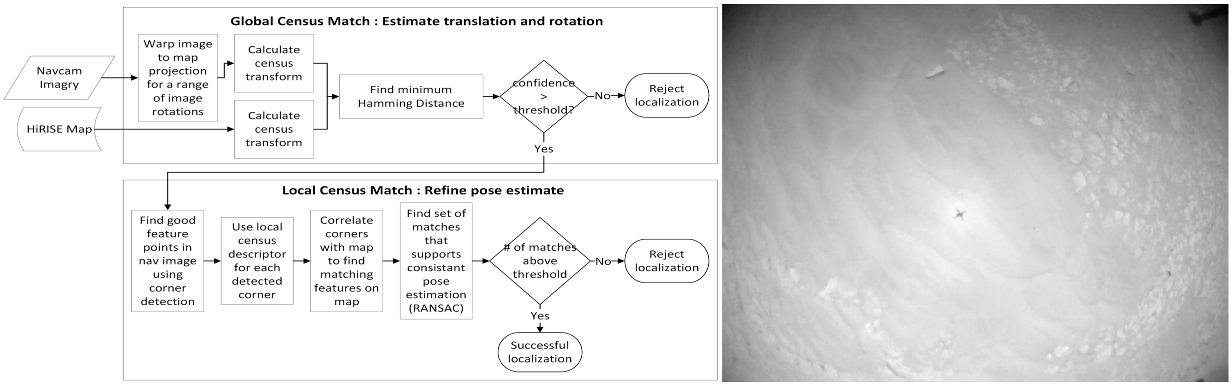

Figure 2: (Left) Process diagram of AGL. (Right) Navcam image from Flight 59.

We divide the localization process into two steps using the census transform [8]. First, a rough estimate is computed using a global census match. Second, the estimate is refined using a local census match. A flow chart of the whole process is shown in Figure 2 (left). As a high level summary, the global census match compares the navcam image to the whole map area. This comparison finds the map location with the highest correlation with the whole image, but lacks granularity in estimating rotational error. The local census match calculates a homography between the image and the map, which means it can capture rotational error very well. Local census matching uses a brute force feature matcher, whose computational complexity is quadratic with the number of feature points. By using the global census match to narrow down the initialization, the local census match can avoid ballooning computational costs.

Results

We were able to improve 23 out of 36 flights' pose estimate using our onboard global localization method. However, we were unable to identify a metric threshold which would allow correct pose estimates to be selected with high precision. In order to deploy this solution, we need to identify solutions which have <10m translation error and <1 degree heading error with 100% precision. This is a very strict criteria, but is needed to ensure the safety of the helicopter.

In addition to challenges selecting a threshold, our method was not able to bound the heading error to 1 degree. Overall, the heading error was improved, with a mean improvement of 1.5 degrees. The training set did identify a threshold of 17 inliers that would bound the heading error to 5 degrees with 100% precision. However no threshold exists that would bound the heading error to 1 degree with 100% precision. If 1 degree is a requirement, different registration methods should be explored.

CONCLUSION

In this paper, we present the onboard framework for automated health assessment and global localization designed for the Mars Helicopter, Ingenuity. Although Ingenuity's mission is over, we hope that this work will help to enable multi-flight operations on future space missions such as Mars Science Helicopter [9], Dragonfly [10], and Mars Inspection Drone [11].

Acknowledgements

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

References and Notes

1. Anderson JL, Karras JT, Cacan M, Kubiak G, Pyrzak G, Dor H, Pipenberg B. Ingenuity, one year of flying on mars. In2023 IEEE Aerospace Conference 2023 Mar 4 (pp. 1-18). IEEE.
2. Grip HF, Conway D, Lam J, Williams N, Golombek MP, Brockers R, Mischna M, Cacan MR. Flying a helicopter on Mars: How ingenuity's flights were planned, executed, and analyzed. In2022 IEEE Aerospace Conference (AERO) 2022 Mar 5 (pp. 1-17). IEEE.
3. Tipaldi M, Bruenjes B. Spacecraft health monitoring and management systems. In2014 IEEE Metrology for Aerospace (MetroAeroSpace) 2014 May 29 (pp. 68-72). IEEE.
4. Kirschenbaum L. A Generic Spacecraft FDIR System. In2021 IEEE Aerospace Conference (50100) 2021 Mar 6 (pp. 1-7). IEEE.
5. Charbuty, Bahzad, and Adnan Abdulazeez. "Classification based on decision tree algorithm for machine learning." *Journal of Applied Science and Technology Trends* 2, no. 01 (2021): 20-28.
6. *Mars 2020 Terrain Relative Navigation HiRISE Orthorectified Image Mosaic*, USGS Astrogeology Science Center, 24 July 2020, doi.org/10.5066/P9QJDP48.
7. Scaramuzza D, Fraundorfer F. Visual odometry [tutorial]. IEEE robotics & automation magazine. 2011 Dec 8;18(4):80-92.
8. Zabih, Ramin, and John Woodfill. "Non-parametric local transforms for computing visual correspondence." *Computer Vision—ECCV'94: Third European Conference on Computer Vision Stockholm, Sweden, May 2–6 1994 Proceedings, Volume II 3*. Springer Berlin Heidelberg, 1994.
9. Tzanetos T, Bapst J, Kubiak G, Tosi LP, Sirlin S, Brockers R, Delaune J, Grip HF, Matthies L, Balaram J, Withrow-Maser S. Future of Mars rotorcraft-Mars science helicopter. In 2022 IEEE Aerospace Conference (AERO) 2022 Mar 5 (pp. 1-16). IEEE.
10. Lorenz RD, MacKenzie SM, Neish CD, Le Gall A, Turtle EP, Barnes JW, Trainer MG, Werynski A, Hedgepeth J, Karkoschka E. Selection and characteristics of the Dragonfly landing site near Selk crater, Titan. The Planetary Science Journal. 2021 Feb 8;2(1):24.
11. Bețco D, Pârvu PV, Ciudin S. Autonomous navigation for the Martian Inspection Drone. Acta Astronautica. 2024 Mar 1;216:55-63.