

IAC-16.D1.2.5x33862

Self-supervised learning as an enabling technology for future space exploration robots: ISS experiments

Kevin van Hecke^{*a}, Guido C.H.E. de Croon^{*a}, Daniel Hennes^b,
Timothy P. Setterfield^c, Alvar Saenz-Otero^c, and Dario Izzo^b

^a MAV-lab, Control and Simulation, Faculty of Aerospace Engineering, TUDelft, the Netherlands

^b Advanced Concepts Team of the European Space Agency, ESTEC, Noordwijk, the Netherlands

^c Space Systems Lab, Massachusetts Institute of Technology, Cambridge, MA, USA

* Corresponding Authors: k.g.vanhecke@tudelft.nl, g.c.h.e.decroon@tudelft.nl

Abstract

Although machine learning holds an enormous promise for autonomous space robots, it is currently not employed because of the inherent uncertain outcome of learning processes. In this article, we investigate a learning mechanism, Self-Supervised Learning (SSL), which is very reliable and hence an important candidate for real-world deployment even on safety-critical systems such as space robots. We introduce a novel SSL setup that allows a stereo vision equipped robot to cope with the failure of one of its cameras. We present preliminary results from an experiment on the International Space Station (ISS) performed with the MIT/NASA SPHERES VERTIGO satellite. The presented experiments were performed on October 8th 2015 on board the ISS. The main goals were (1) data gathering, and (2) navigation on the basis of stereo vision. First the astronaut Kimiya Yui moved the satellite around the Japanese Experiment Module to gather stereo vision data for learning. Subsequently, the satellite freely explored the space in the module based on its (trusted) stereo vision system and a pre-programmed exploration behavior, while simultaneously performing the self-supervised learning on board. The two main goals were successfully achieved, representing the first online learning robotic experiments in space. These results lay the groundwork for a follow-up experiment in which the satellite will use the learned single-camera distance estimation for autonomous exploration in the ISS, and are an advancement towards future space robots that continuously improve their navigation capabilities over time, even in harsh and completely unknown space environments.

Keywords: Persistent self-supervised learning, stereo vision, monocular depth estimation, space robotics.

Acronyms/Abbreviations

SSL	Self-Supervised Learning
JEM	Japanese Experimentation Module
SPHERES	Synchronized Position Hold Engage and Reorient Experimental Satellite
VERTIGO	Visual Estimation for Relative Tracking and Inspection of Generic Objects
ISS	International Space Station
ROC curve	Receiver Operating Characteristic curve

relevant for planetary explorers that are by definition sent to areas of which little is known. Despite these advantages, machine learning is not yet used in space robotics; the major reason for this is that it introduces extra risk for missions with extremely small margin for error.

A potential suitable candidate for learning on space robots is Self-Supervised Learning (SSL), since it is a reliable learning method that allows robots to adapt to their environment. In an SSL setup, a robot extends its capabilities by using a trusted, primary sensor cue to train a still unknown secondary sensor cue. The most well-known example of SSL is its use on autonomous driving cars such as Stanley [6-10] for the recognition of drivable terrain. Stanley used a laser scanner to detect drivable terrain close by. It then used the outputs from the laser scanner as classification targets for a supervised learning process. This process learned a function that maps colors in the camera image to the classes of "drivable" and "non-drivable". Since the camera has a much longer range than the laser scanner, Stanley could see where the road was going for much further distances, allowing it to move more quickly and win the grand DARPA challenge. However, SSL can be

1. Introduction

Future space missions will increasingly rely on the help of robotic systems and in some cases even be performed purely by fully autonomous robots. Currently, robots are either tele-operated remotely [1-3] or perform autonomous tasks such as driving short stretches in a pre-programmed manner [4-5]. Hence, robots still lack the ability to learn from their environment.

Learning can be very advantageous, since it decreases pre-deployment engineering effort and allows the robot to adapt to specific properties of its potentially yet unknown environment. This last advantage is especially

more broadly applied to other modalities and for other purposes [11-13].

We recently introduced a novel SSL setup that allows a robot equipped with stereo vision to cope with a potential failure of one of its cameras [14]. The idea is that the robot learns how to see distances in a single, still image while operating in its environment. To this end, it will learn a function that maps textures in an image to the distances obtained with its stereo vision system. As opposed to almost all previous work on SSL, our proposed setup will require the robot to learn a model that persists over time, which is why we term it *persistent SSL*. We have previously successfully tested persistent SSL for the navigation of autonomous drones on Earth [14], which showed its potential in terms of reliability for application in space.

In this article, we will present preliminary results from an experiment on the International Space Station (ISS) performed with the MIT/NASA SPHERES VERTIGO satellite, which is equipped with a stereo vision system that allows it to perceive depth and navigate by itself. The presented preliminary experiments, prepared by a mixed team from TU Delft, the European Space Agency (ESA) and the Massachusetts Institute of Technology (MIT), were performed on October 8th 2015 on board the ISS. The main goals were (1) data gathering, and (2) navigation on the basis of stereo vision. Going beyond the goals of the experiments, the SPHERES VERTIGO satellite also performed the learning during operation, making it – to the best of the authors' knowledge – the first learning robot in space.

The remainder of the article is structured as follows. In Section 2 we explain the materials and methods used in the experiment. Subsequently, in Section 3 we discuss the experimental results, which are discussed in Section 4. Finally, we draw conclusions in Section 5.

2. Material and methods

In recent work, we proposed a persistent Self Supervised Learning (SSL) method to solve the proof of concept problem of learning monocular distance estimation on board a robot [14]. Persistent SSL is able to use on board available training data to train another algorithm online, which may use a different sensor. Due to the fact that ground truth is available during most of the run-time, performance guarantees can be made on the outcome of the learned estimate. Moreover, the learning algorithm is trained in exactly the same environment as the deployment environment, while providing large amounts of training data (since the labelling is done automatically by a trusted algorithm based on a trusted sensor cue available on the drone). Advantages include possible extrapolation to results better than the trusted cue (since the cue on which the learned algorithm was trained may be better in some situations), while possibly being less CPU intensive,

and most importantly for this work, providing another adaptive and redundant cue to become more resistant against sensor failure etc.

In order to perform the experiment, we utilized the SPHERES test bed system, which was launched to the ISS in 2006. This system is meant as a test bed platform, to give engineers a chance to test algorithms on real spacecraft in a microgravity environment before deploying them on multi-million dollar satellites. A stereo vision upgrade (VERTIGO) was installed in 2013. Figure 1 shows a satellite of the SPHERES platform with the VERTIGO upgrade. A SPHERES satellite contains 12 CO₂ thrusters to provide full 6-DOF control in the micro gravity environment of the ISS. The VERTIGO system consists out of a monochrome 640x480 resolution stereo camera, and a 1GHz Via x86 embedded computer running Ubuntu. Using this platform, we tested our persistent-supervised learning method on board the ISS, inside the Japanese Experiment Module (JEM).

2.1 Experiments

We have conducted two experiments. The first experiment was a short experiment (2 minutes). The autonomous behaviour and tuned control algorithms required ground truth data since the persistent-SSL algorithm had not been previously tested in the ISS; the first experiment collected this ground truth data and served as a contingency data collection experiment. The satellite was manually moved through the JEM volume by the astronaut in a way that simulated the autonomous behaviour as closely as possible, keeping the untested 6-DOF parts out of the loop. The astronaut was provided haptic feedback to move the satellite differently by means of short bursts of the thrusters. A manual and quick briefing before the test was given to the astronaut on how to move the satellite.

The second, autonomous, test was set to take 10 minutes. Figure 2 shows the system diagram. The satellite starts by using its stereo vision to explore the SPHERES-volume inside the JEM, by moving in a straight line until an obstacle (JEM wall) is detected. After detection, the satellite rotates in a random direction, checks if there is no obstacle in that direction, and then proceeds in another straight line towards that direction. In the meantime, persistent-SSL is training the estimation of a single (average) depth based on the stereo vision depth information and the monocular image data from one of the cameras of the stereo vision system. After 7 minutes, or if the algorithm achieved a low error rate on its learned results, the monocular vision was set to be enabled in the behaviour loop. The error rate is determined by means of a ROC curve analysis, like in our previous work on the small quadcopter [14]. Since the SPHERES volume consists of two open ends, the stereo vision will not detect an

obstacle in those regions, but the satellite is not allowed to venture of these regions. In those cases, the SPHERES global metrology (sonar based pseudo-GPS) is used to determine if the satellite is within the volume. In case the satellite is leaving the volume, the satellite stops, rotates randomly to a direction inwards to the volume, and starts a straight line towards that new direction. A safety override control manoeuvre is activated when the satellite moves close to the outer ranges of the expected range of the global metrology system, in which case attitude control of the satellite is disabled and full thruster power is used for position control to return to the test volume.




Figure 1 - SPHERES VERTIGO

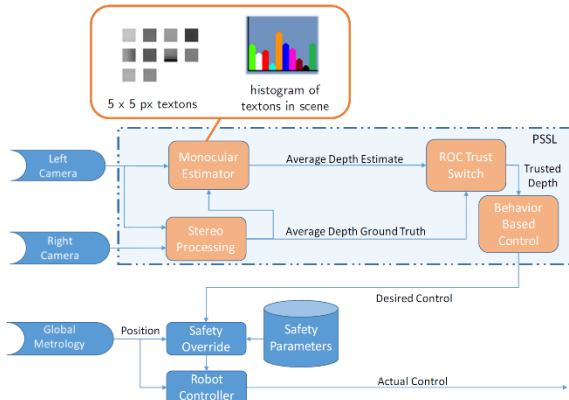


Figure 2 - System overview

3. Results

SPHERES Test Session 74B was conducted on October 8th, 2015, with astronaut Kimiya. The test session contained six tests; this paper describes the results of test four and test five in particular. Both tests four and five were done twice due hardware failures. A summary of the results of these tests is given in Table 1. Videos of each test and their raw data can be viewed online at https://www.youtube.com/playlist?list=PL_KSX9GOn2_P-NpLU8DIS-PmgqlfB3a_-A

Table 1. Test results overview.

Test	Duration	Images*	Notes
T4.1	~50s	596/144	IR overload reset
T4.2	120s	1131/267	Success

T5.1	~30s	261/123	Battery empty
T5.2	210s	1293/605	IR overload reset

*Acquired / processed and learned on-board

Manual tests #1 (T4.1 & T4.2)

The objective of the manual test, was for the astronaut to move the SPHERES satellite through the volume, in order to obtain images of the surroundings. Kimiya picked up the instructions well and attempted to execute some useful trajectories. However, after 60 seconds, a satellite reset occurred. Offline analysis shows the reason for the reset as a disturbance on the global metrology infrared detector on the satellite, which caused a high priority interrupt to continuously fire, which caused a system reset by means of a watchdog.

The 3D path of T4.1, as undertaken by Kimiya before the reset, is depicted in Figure 3.

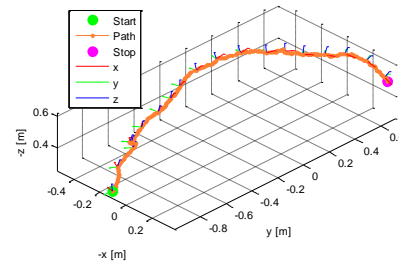


Figure 3 - Satellite flight path of T4.1

In the time before the unexpected reset, 596 stereo images have been acquired of which 144 were directly used for training.

Due to the reset in T4.1, a second run of T4 was attempted and successfully finished. The learned data from T4.1 was automatically concatenated to T4.2. During T4.2 an additional 1131 images were acquired, of which 267 were directly used for online training. A 3D plot of the trajectory of the satellite is shown in Figure 4.

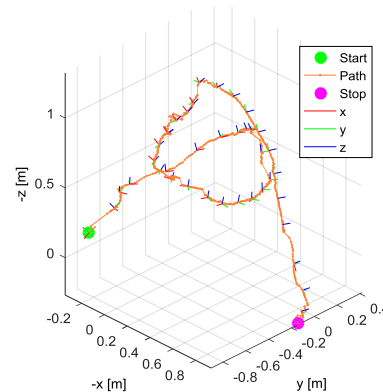


Figure 4 - Satellite flight path of T4.2

The stereo images acquired during T4 were analysed in real-time and as ground truth the average disparity for each stereo pair was calculated. The results of this can be seen in Figure 5.

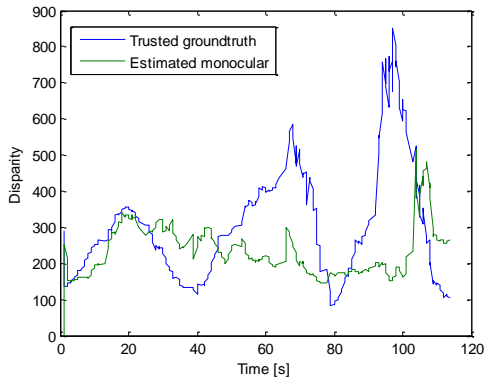


Figure 5 - Average disparity during T4.2

The estimated disparity is based on the self-supervised learned algorithm trained on the 144 images from T4.1 during T4.2. As can be seen, some correlation is visible, but the training set was too small to properly estimate the disparity at this point.

Autonomous tests (T5.1 & T5.2)

The autonomous test was also attempted twice. The first time, T5.1, a battery empty failure occurred, which caused the VERTIGO system to automatically shut down after 30 seconds, and loss of control happened after 11 seconds. The 3D path is depicted in Figure 6. During its first and only trajectory, the satellite managed to evade another stationary satellite (that was placed there mistakenly) using its trusted stereo algorithm. This evasion maneuver is shown as the first green circle in the plot. The subsequent vision turns were commanded by the algorithm, but not executed due to the battery problem. During T5.1 an additional 261 images were acquired, of which 123 were analyzed directly for training.

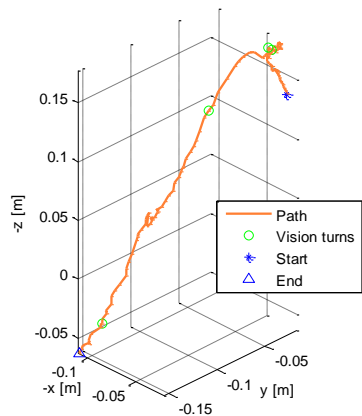


Figure 6 - Satellite flight path of T5.1

As T5.1 failed due to power loss, another attempt was done. Unfortunately, the second attempt failed after 210s. Offline analysis showed the cause to again be the IR disturbance problem that also ended T4.1. Before that time, the satellite was able to make several crossings through the volume exploring and learning simultaneously. Also, another 1293 images were acquired of which 605 analyzed during the test. The 3D flight path of T5.2 is depicted in Figure 7.

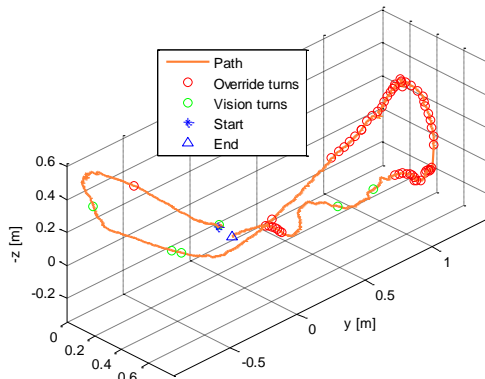



Figure 7 - Satellite flight path of T5.2

The first turn was initiated by the global metrology based safety override, as the satellite ventured out of the open end of the volume. The safety override has therefore been proven to work as expected. The second and third turn were both stereo vision avoidance commands, avoiding collision with the JEM wall. The fourth turn was caused by the safety override, but happened at too high a speed. This caused satellite to drift too far out of the allowed volume causing the emergency safety override to become active. The satellite diverted all control power to position control, disabling attitude control, in order to move back into the volume as quickly as possible. In a best case scenario, this would not have been necessary. For future missions we plan to reduce the top speed of the satellite such that




Figure 8 - Results from T5.2 overview

this will not happen again. The last two turns are caused by stereo avoidance commands, again attempting to avoid the JEM wall. However, before the satellite had a chance to execute the avoidance manoeuvre, the IR reset failure interfered.

Other, previous experiments on the SPHERES VERTIGO were much less affected than ours by the IR reset failure. The current hypothesis is that it has to do with the more complete exploration as performed by the autonomous stereo vision based behaviour. This brings the satellite closer to interfering electronics in the module. Various solutions are under evaluation to prevent the problem in future experiments.

A video of the results calculated on the satellite during the experiment in the ISS can be viewed online: https://www.youtube.com/watch?v=6mk0om2TPTo&index=8&list=PL_KSX9GOn2P-NpLU8DIS-PmgqlfB3a_-A

A still image from this video can be seen in Figure 8, this video was also visible for Kimiya during the tests by means of a Wi-Fi stream. The image contains the raw input images from the left camera (top left in the figure) and right camera (top, middle), the disparity map (top right – red is close, blue is far), and a graph of the average disparity over time (bottom time line plot). The average disparity from the trusted stereo vision algorithm is shown in blue, the learned estimates are shown in green, while predictions of the robot of previously unseen cases are shown in red. When the trusted disparity supersedes the threshold, an avoidance manoeuvre is commanded.

We zoom in on two moments in time during the learning on board the satellite. Figure 9 and Figure

10 show the learned disparities (green) and predicted disparities (red) just before $t = 141$ seconds and after $t = 141$ seconds. The predictions before $t = 141$ seconds do not seem to correlate much with the "ground-truth" stereo vision estimates. However, after learning on these samples, the predictions on new, unseen samples do correlate well with the ground-truth. The predictions can be evaluated objectively with respect to the stereo vision threshold – resulting in a classification problem setting. Before learning at $t = 141$ seconds, the True Positive Rate (TPR) of the predictions is 0.2 while the False Positive Rate (FPR) is 0.6. After learning that part, the predictions are indeed objectively better, with a TPR of 0.7 and FPR of 0.3. These results correspond to results obtained on earth with fewer degrees of freedom [14]. They show that the learning is successful, but that the amount of gathered data is not yet enough to cover the entire environment.




Figure 9 - Estimated disparity before learning the particular scene at $t=141$ s.




Figure 10 - Estimated disparity at $t = 212s$, after learning the particular scene up to $t=141s$.

4. Conclusions

We have presented preliminary results from a Self-Supervised Learning (SSL) experiment on the International Space Station (ISS) performed with the MIT/NASA SPHERES VERTIGO satellite. The main goals of the experiment were (1) data gathering, and (2) navigation on the basis of stereo vision. Both goals were successfully achieved, although the experiments were hampered by automatic resets triggered by an interference of the IR detector of the SPHERES satellite. During both parts of the experiment, the satellite was learning online to map the appearance of the environment to the distance estimates from its stereo vision system. Despite the extremely limited training time, some successful generalization of the learned mapping to unseen images can be observed.

These first robotic learning experiments in space hold a promise for follow-up experiments in which the satellite will use the learned mapping to navigate with only a single camera. For a follow-up experiment, the following main insight from the current experiment should be taken into account. The combination of a limited experiment time with a relatively slow speed of the robot satellite implies that it will be difficult to learn the entire available space in the module. This is especially true given the fact that the robot in space can move with 6 DOF, leading to a high variety in environment appearance. Given a limited experiment time, the movement space of the satellite should also be more limited, e.g., by not having it travel to the open ends of the module. As a by-effect, this will avoid the problems we experienced when our emergency override system disabled attitude control to enter back into the central space. This will hopefully allow learning and switching to monocular vision control in a single test session.

References

[1] Heikkilä S.S., Halme A. and Schiele A., (2012), 'Affordance-based indirect task communication for astronaut-robot cooperation', *Journal of Field Robotics*, Vol. 29, Issue 4, pp. 576 – 600
[2] Hirzinger, G., Brunner, B., Dietrich, J., & Heindl, J. (1993). Sensor-based space robotics-ROTEX and its

teleroptic features. *IEEE Transactions on robotics and automation*, 9(5), 649-663.

[3] Hirzinger, G., Landzettel, K., Brunner, B., Fischer, M., Preusche, C., Reintsema, D., ... & Steinmetz, B. M. (2004). DLR's robotics technologies for on-orbit servicing. *Advanced Robotics*, 18(2), 139-174.

[4] Matthies, Larry, et al. "Computer vision on Mars." *International Journal of Computer Vision* 75.1 (2007): 67-92.

[5] Maimone, M. W., Leger, P. C., & Biesiadecki, J. J. (2007, April). Overview of the mars exploration rovers' autonomous mobility and vision capabilities. In *IEEE international conference on robotics and automation (ICRA) space robotics workshop*.

[6] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann et al., "Stanley: The robot that won the darpa grand challenge," in *The 2005 DARPA Grand Challenge*. Springer, 2007, pp. 1–43.

[7] D. Lieb, A. Lookingbill, and S. Thrun, "Adaptive road following using self-supervised learning and reverse optical flow." in *Robotics: Science and Systems*, 2005, pp. 273–280.

[8] A. Lookingbill, J. Rogers, D. Lieb, J. Curry, and S. Thrun, "Reverse optical flow for self-supervised adaptive autonomous robot navigation," *International Journal of Computer Vision*, vol. 74, no. 3, pp. 287–302, 2007.

[9] R. Hadsell, P. Sermanet, J. Ben, A. Erkan, M. Scoffier, K. Kavukcuoglu, U. Muller, and Y. LeCun, "Learning long-range vision for autonomous off-road driving," *Journal of Field Robotics*, vol. 26, no. 2, pp. 120–144, 2009. [Online].

[10] U. A. Muller, L. D. Jackel, Y. LeCun, and B. Flepp, "Real-time adaptive off-road vehicle navigation and terrain classification," in *SPIE Defense Security, and Sensing. International Society for Optics and Photonics*, 2013, pp. 87 410A–87 410A.

[11] J. Baleia, P. Santana, and J. Barata, "On exploiting haptic cues for self-supervised learning of depth-based robot navigation affordances," *Journal of Intelligent & Robotic Systems*, pp. 1–20, 2015.

[12] H. Ho, C. De Wagter, B. Remes, and G. de Croon, "Optical flow for self-supervised learning of obstacle appearance," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on. IEEE*, 2015, pp. 3098–3104

[13] Otsu, K., Ono, M., Fuchs, T. J., Baldwin, I., & Kubota, T. (2016). Autonomous Terrain Classification With Co-and Self-Training Approach. *IEEE Robotics and Automation Letters*, 1(2), 814-819.

[14] van Hecke, K., de Croon, G., van der Maaten, L., Hennes, D., & Izzo, D. (2016). Persistent self-supervised learning principle: from stereo to monocular vision for obstacle avoidance. *arXiv preprint arXiv:1603.080*

