# Onboard classification to guide capture downlink using the HYPSO-1 satellite

**Conference Paper** · October 2024

CITATIONS
0

READS
8

**6 authors**, including:

Simen Berg
Norwegian University of Science and Technology
**11** PUBLICATIONS **46** CITATIONS
SEE PROFILE

Dennis Langer
Norwegian University of Science and Technology
**25** PUBLICATIONS **94** CITATIONS
SEE PROFILE

Corrado Chiatante
Norwegian University of Science and Technology
**6** PUBLICATIONS **3** CITATIONS
SEE PROFILE

Roger Birkeland
Norwegian University of Science and Technology
**74** PUBLICATIONS **381** CITATIONS
SEE PROFILE

Post Print
75th International Astronautical Congress (IAC), Milan, Italy, 14-18 October 2024.


IAC–24–B.4.3.6x85269

# Onboard classification to guide capture downlink using the HYPSO-1 satellite

**Simen Berg[a*], Dennis D. Langer[b], Corrado Chiatante[c], Roger Birkeland[c], Jon A. Justo[d], Tor A. Johansen[d]**

[a] *Department of Engineering Cybernetics, Norwegian University of Science and Technology, O.S. Bragstads Plass 2D, 7034 Trondheim, Norway, simen.berg@ntnu.no*
[b] *Department of Marine Technology, Norwegian University of Science and Technology, Trondheim, Norway,*
[c] *Department of Electronic Systems, Norwegian University of Science and Technology, Trondheim, Norway,*
[d] *Department of Engineering Cybernetics, Norwegian University of Science and Technology, Trondheim, Norway,*
[*] *Corresponding author*

**Abstract**

The HYPSO-1 6U CubeSat is equipped with a hyperspectral imager and has acquired and downlinked over 2000 hyperspectral images during its first 2.5 years on orbit. The payload can acquire many images per day, but the satellite downlink and available ground stations imposes a capacity limit of about 5-6 captures per day. Since the launch, the team has operated HYPSO-1 and continued updating it with functionality, including an onboard classification algorithm currently applied to the acquired images. The onboard classification is a Support Vector Machine (SVM) utilising a binary decision tree that is labeling pixels into 11 different classes, including clouds, land, and water. Until now, these labels have been exclusively analysed after downlinking to a ground station. In this research, we perform onboard analysis of the labeled images to generate more knowledge onboard the satellite. The knowledge gained is intended to be used to prioritise the downlink of captures. The main idea is to increase the satellite's coverage by acquiring more images than the satellite can downlink, and utilise onboard classification to filter out acquisitions of poor quality or relevance. We explore the viability of this idea by extending the onboard classification model to output the percentage of pixels per class, and showing the necessary architectural changes to the scheduling software and onboard software to implement the functionality. The result shows that the concept is viable but requires improving the classifier to get reliable classification.

**Acronyms/Abbreviations**
Support Vector Machine (SVM)
Convolutional Neural Network (CNN)
Onboard Processing Unit (OPU)

## 1. Introduction

CubeSats have been used in educational satellite missions due to their lower deployment cost compared to traditional, monolithic satellites, while still proving valuable in targeted missions [1]. Similar to the Copernicus program's expansion missions, CubeSats can also play a key complementary role in supporting large observation programs by filling unaddressed niches [1]. However, their small size limits critical resources like power and sensor sizes, leading to a reduced capacity compared to conventional satellites. Such constraints have been addressed by utilising multiple small satellites in large constellations, such as Planet Labs' constellation [2]. With a large amount of satellites, they can provide global coverage and are better equipped for low-latency coverage of specific target areas. However, such coverage cannot be provided with a single CubeSat.

An example of educational satellite missions is HYPSO (Hyperspectral Small Satellite for Ocean Observation). The first CubeSat, HYPSO-1, was launched on the 13th of January 2022. The satellite is equipped with a hyperspectral imager and has, since its launch, successfully acquired and fully downlinked over 2000 hyperspectral cubes, with the main mission being focused on marine observations, especially harmful algae blooms. The HYPSO-1 satellite can downlink around 5-6 captures per day [3], but accounting for energy and onboard storage capacity, it can acquire significantly more captures daily. The main limitation is onboard- and downlink data transfer rates. In Fig. 1, we illustrate the image downlink procedure. The payload first acquires the data, which is then transferred - referred to as *buffering* - to the payload controller for temporary storage before it is downlinked to the ground station via the S-band radio. The payload controller is used as temporary storage since it has a much faster communication interface to the S-band radio than the payload has to the S-band radio. The payload only has a CAN bus interface to the S-Band radio, whilst the payload controller has a higher data rate UART interface. As seen in Fig. 1, the CAN bus

has a significantly lower throughput than the S-band and would become the limiting factor if images were directly downlinked from the payload.

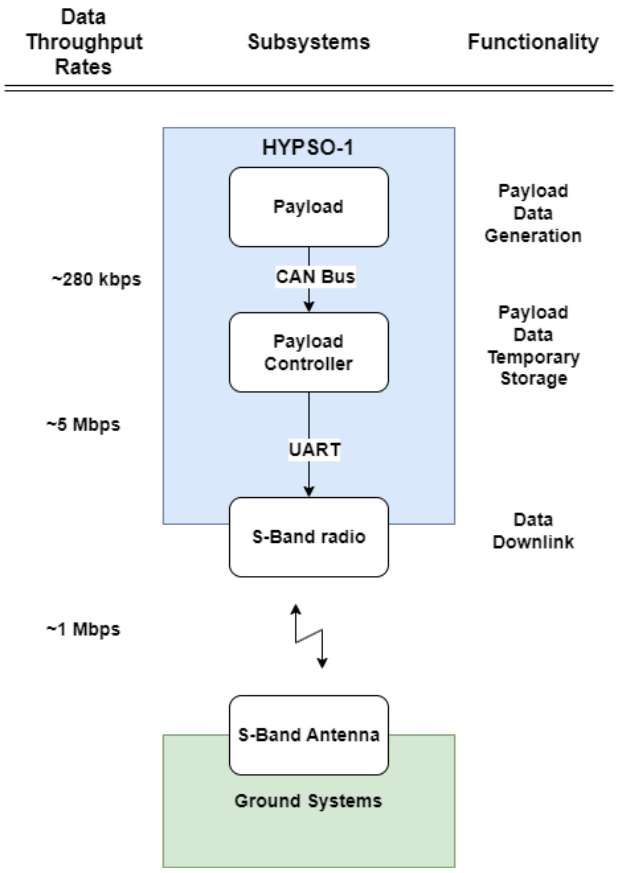


Fig. 1. Procedure for performing imaging downlink.

Due to the satellite's limited downlink capacity, the coverage is also limited. However, its strength lies in its agility, meaning its ability to control its attitude to acquire images off-nadir. The satellite performs a small selection of targeted captures per day to monitor specific areas, with most images being acquired pointing sideways off-nadir. The areas to image are finally selected by an in-house automated system from a list of desired targets as described in [4]. There is a possibility of human intervention if desired. When scheduling the imaging acquisitions, cloud predictions from the Norwegian Meteorological Institute are used [5], and planning is performed at least three times per week to have up-to-date weather reports. However, not all captures are successful. Pointing anomalies, inaccurate cloud predictions, space weather justifying powering off non-essential subsystems as a precaution, unexpected reboots or downlink errors are some causes of failures. With data throughput as a main constraint, the satellite is not as suitable for detecting new phenomena as it is for monitoring existing phenomena since the amount of daily images is limited. Thus, the capacity of the satellite is better used for monitoring the development of known algae blooms. However, since our imaging capability is greater than our downlink capability, we want to increase our daily image acquisitions and utilise onboard processing to only downlink the data of highest relevance and quality. In this paper, we present a novel approach to enhance our operations by using onboard classification to guide our imaging downlink and discuss what architectural changes to our payload are required to implement it.

## 2. Cloud Detection & Onboard Classification in the HYPSO project

Onboard cloud detection and thus download prioritisation has become an objective for multiple optical satellite owners and operators in recent years, such as in ESA's Copernicus Hyperspectral Imaging Mission for the Environment (CHIME), Φ-sat-1, Φ-sat-2 missions as well as commercial actors such as Kuva Space [6–9]. These missions aim to make cloud masks onboard and use the new knowledge to apply lossy compression to the cloudy pixels to spare downlink resources. From [8], it is stated that clouds cover >54% of the Earth's land surface and 68% of the oceans. Performing lossy compression of pixels affected by clouds has the potential to significantly reduce resources misused on downlinking unusable data.

Within the HYPSO mission, we started working towards onboard processing to guide downlink at an early stage of the orbital lifetime of HYPSO-1. In [10], Birkeland et al. explored the idea of utilising the compression ratio of the hyperspectral images onboard the satellite to get some information about the general quality of the capture. It was never taken into practical use, but we have since developed classification models. One is a Deep Learning Convolutional Neural Network (CNN) called *1D-Justo-LiuNet*, which was trained in [11], and the other model successfully run in-orbit is a Machine Learning Support Vector Machine (SVM) utilising a binary decision tree developed by Røysland in [12, 13]. On the one hand, the CNN classifies each pixel in the hyperspectral images into one of the three classes *land, water, clouds*. Model inference runs on raw data without calibration or correction to minimise computational complexity, as previous results show comparable results between inference on raw data and radiance [11]. On the other hand, the SVM classifies the images into the 11 different classes *water, strange water, light forest, dark forest, urban, rock, ice, sand, thick clouds, thin clouds*, and *shadows*. The

SVM model performs radiometric calibration, described in [12, 13], before running classification inference. The class *strange water* is essentially a category consisting of water with some distortion, e.g., from sediments or algae. Although the CNN offers significantly higher accuracy, the SVM, simpler and with shorter inference time, classifies each pixel into one of many more categories, yet at the cost of reduced accuracy. An example image is shown in Fig. 2. The image marked *B* is the result of the SVM and shows it has misclassified a large portion of the land as the label *sand* when it should probably be *dark forest*.

The decoded classified images are henceforth called *labeled images*. In the decoded labeled image by the CNN, orange denotes *land*, gray indicates *clouds*, and blue shows *water*. For the classified image by the SVM, the legend of the decoded SVM is shown in Fig. 3 and is obtained from [13].

The satellite's hyperspectral imager has a spatial resolution of approximately 120 x 630 m/pixel, dependent on the acquisition angles, mode and camera setting [14, 15], and typically acquires hyperspectral cubes with 598 x 1092 pixels in the spatial domain. Thus, the coverage within a single capture is in the region of tens of thousands of square kilometres, dependent on the off-nadir angle of the acquisition.

Compared to the missions with the goal of running onboard cloud detection mentioned earlier, the HYPSO project has a slightly different problem statement. We have a satellite capable of capturing many more images than we can downlink per day. The images we acquire are targeted to cover a specific area and typically cover a much larger region than the area of interest. This does not mean that the additional pixels are not of interest, but the main area of interest is often contained within a small portion of the full image acquisition. A capture can be highly successful even though, for example, 80% of the capture is affected by clouds, or a total failure even if 5% is affected by clouds, depending on whether the clouds cover the specific target area or not. Therefore, we want to utilise onboard classification to better understand whether our specific imaging target was successfully captured.

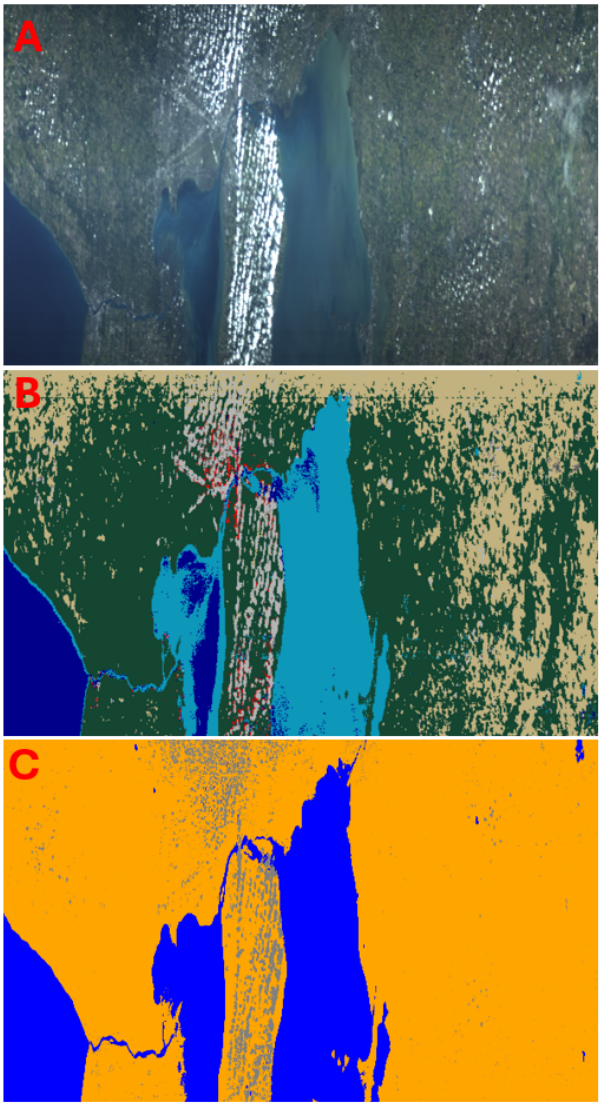


Fig. 2. Comparison of the classification models. The image was acquired by the HYPSO-1 satellite on the 18 th of September 2024 and is of Lake Erie on the border between the USA and Canada. The picture marked *A* is an RGB composite of the hyperspectral cube generated on ground. The picture marked *B* is the labeled image using the SVM run in-orbit. The picture marked *C* is the labeled image using the CNN run on ground. Figure C is decoded using the software published by Justo et al. in [11].

## 3. Selection of classification model and approach

To enhance our satellite operations, we can choose different approaches. As described in Section 2, we want to determine if we got successful imaging of the specific target area, not necessarily the entire capture area.


Fig. 3. Color-mapping of the decoded SVM classes. Figure obtained from [13]

*3.1 Onboard georeferencing*

One of the approaches considered was performing onboard direct- and indirect georeferencing to find the target area within the acquired image. This is further described in Chiatante et al. in [16], utilizing the CNN from [11] to make cloud, land, and water masks as input to the indirect georeferencing algorithm. It would require maps of the imaged area to be uplinked to the onboard payload computer before imaging. However, the georeferencing cannot be too computationally complex for it to be viable in nominal operations. This approach was disregarded at the current point in time as it was deemed too complex within the scope of this paper. It could become a viable option in the future, but it requires further research.

*3.2 Cloud percentage*

Another approach could be to utilise build upon the previous idea described by Birkeland et al. in [3] of utilising the compression ratio of the hyperspectral images to estimate the cloud cover. One way of improving it is using a histogram plot of intensity values, which could output the percentage of overexposed pixels, assuming that those are due to clouds. As proposed in [11], the same could be done by using the CNN by running it onboard the satellite. Additional information that would be interesting to extract from a classification method would be to check if there is any coastline in the image, which is necessary for georeferencing to be performed on the ground in post-processing. The histogram plot could possibly also be used for mapping coastlines by comparing dark ocean pixels close to brighter land pixels. Still, the logic and testing/training required to set this up essentially means making a new coastline classification model, which is unnecessary in the scope of this work. The use of cloud percentage could be used as a prioritisation between two target areas where one has significantly more clouds than another; a concept described in [11]. However, it would not yield any information about whether the clouds are over the specific target area or not.

*3.3 Target biomes*

The approach selected in this paper was to use the SVM classification model. It is not the most accurate at its current stage, but it labels the captures into a large set of different classes and runs in less than ten seconds on the payload [13]. The goal was to mark the areas with the biome we are interested in and set a threshold for the percentage of pixels of that biome for the capture to be considered a success. A *biome* is here defined as one or more of the existing classes in the SVM model. In a more general sense, it is the desired classification classes for a specific satellite mission. This approach does not give any exact result on whether or not the particular target area itself was cloudless or if the correct place was imaged. Still, it provides more information than only setting all the land masses as the label *land*. If a large pointing error happens during imaging, the labels will most likely not match the target biome. In the HYPSO mission - being focused on marine observations, specifically harmful algal blooms - the biome *strange water* could be utilised to enhance HYPSO-1's ability to **detect** new algal blooms. Accounting for its downlink limitations, the satellite can only acquire images of a small set of target areas daily, where most of the capacity is allocated for monitoring known algal blooms. Thus, the remaining capacity to detect new algal blooms is scarce. However, by utilising onboard classification to guide the image downlink, the satellite could image a larger variety of water bodies. The threshold of the class *strange water* could be set slightly higher than the nominal average within that area, thus triggering an image downlink only if the payload notices signs that could mean the start of a new algal bloom. The biome *strange water* could also be utilised to detect and monitor other phenomena where the water's spectral response is affected by, e.g., sediments, such as meltwater or flooding.

## 4. Software architecture changes - onboard and scheduling

Currently, the planning and scheduling of captures, buffering sequencing to the payload controller, and downlink to the ground are defined in software on the ground. The payload comprises the hyperspectral imager and a device for controlling the imager, called the Onboard

Processing Unit (OPU). At the time of writing, all the scheduled captures are queued for downlinking since the scheduling software is run on the ground with no knowledge of the capture quality. The payload controller - being used for temporary storage of images before being downlinked to the ground - can only store two captures at a time. The scheduling software estimates at what time these captures are expected to be finished downlinking. It then schedules a new image to be buffered from the payload to the payload controller once the previous capture is estimated to be completely downlinked, as shown in Fig. 4.

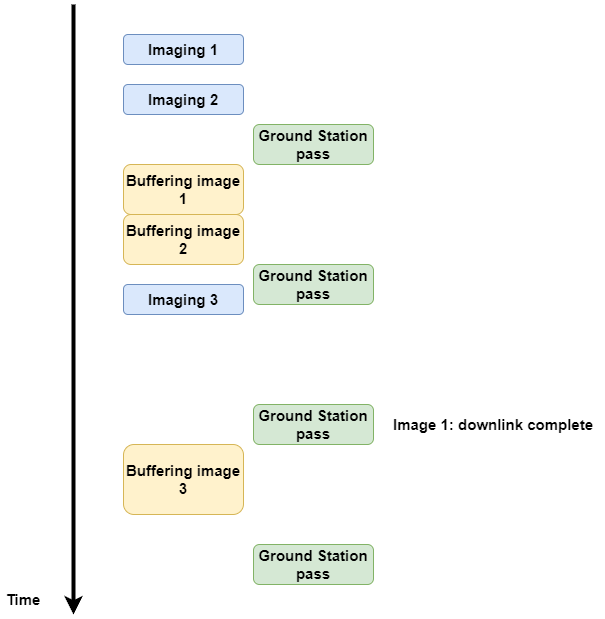


Fig. 4. Example of a downlink schedule estimated by the scheduling software. *Image 3* cannot be buffered to the payload controller until *image 1* is fully downlinked.

The concept of guiding image downlink based on onboard classification is not seen as the novelty in this work as it has been proposed earlier, e.g., in [17], the novelty is the implementation on the HYPSO-1 satellite using the SVM. The goal is for the payload to select which capture to buffer based on the classification output and user-set thresholds. However, we cannot totally disregard all captures not selected for downlink, since the classification output could be inaccurate, their content could hide underlying issues with the satellite. Instead, it would be better to use a queuing system such that if there is downtime with no high-priority captures to downlink, captures of a lower priority will be downlinked. As a minimum, all captures should downlink the labels from the SVM to get a preview of the capture. The information can be used to determine what caused the capture to fail the stated threshold, whether it was a pointing anomaly, payload issue, or inaccurate cloud or biome prediction. Fig. 5 shows how scheduling could be set up with the implementation of a prioritised buffering queue. The satellite would image more than it is capable of downlinking, and then only downlink the captures that are predicted to be the most likely to contain specific targets of interest.

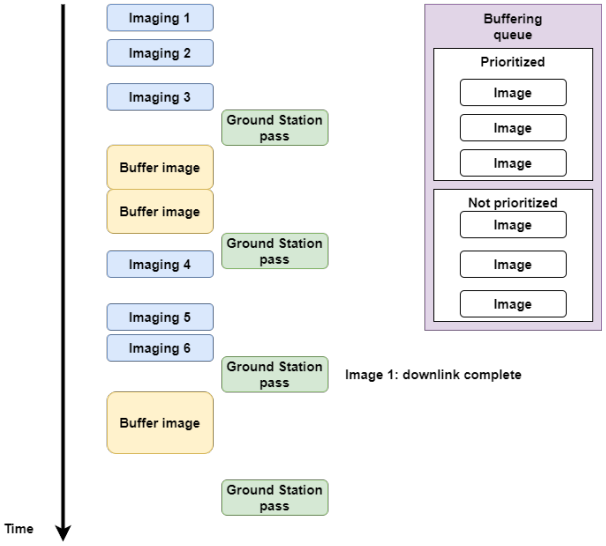


Fig. 5. Example of how scheduling would be set up with the addition of a prioritised buffering queue.

Specific thresholds for each capture target would have to be set as the criteria for the capture to be put into the prioritised buffering queue. These thresholds would be put into the scripts made by the scheduling software on the ground and given to the payload through the scripts. The easiest way for the payload to check if the thresholds are met would be to output a text file containing the percentages of each class within the picture as part of the classification process to get a quantified measure of what the image contains. The thresholds should be a certain percentage of a specific class or a combination of classes. The OPU would then check if the specified classes' thresholds are met and put the capture into the prioritised buffering queue.

## 5. Classification analysis

We made a script to decode the labeled image and record the percentages of pixels per class. To test whether the proposed procedure was viable, we ran this during

normal operations for a few months during the summer of 2024 to get multiple sample images from various scenes and compared the pixel percentages on the ground afterwards. Fig. 6, Fig. 7, and Fig. 8 show the results of two captures at three of the locations we had multiple captures from within the period. All classes where both captures had less than 0.2% of pixels within the class were removed to make the plots more readable. Note that the three figures have different amounts of classes plotted in the tables due to this.

Fig. 6 shows a comparison of two captures acquired of Lake Erie. By looking at the percentage of pixels for the different classes, one of the most significant discrepancies is that *Erie1* is denoted to have the largest amount of *strange water* of the two images. Looking at the RGB composite of the hyperspectral cube and the labeled image, we can clearly see that this is incorrect. The land pixels were incorrectly labeled as *strange water* when they probably should have been labeled as *dark forest* pixels. This is an example of when the downlink prioritisation would not work as intended if the target biome was *strange water*.

Fig. 7 shows two captures of Checleset. Both have a large amount of *water* and *strange water*, but we can see a possible misclassification in the water to the right of Vancouver Island. In the RGB composites, it looks like there are some green patterns that could be algae, but it has not been labeled as *strange water*. *Checleset1* has more *strange water*, *dark forest* and *urban* than *Checleset2*. Even though it does not confirm there are coastlines in the capture, it can be interpreted as the image is more likely to contain coastlines, which is necessary to perform indirect georeferencing during post-processing on the ground. For these captures, this main target of interest was monitoring algae blooms in a small area within the captured region, but surrounding areas were considered additional value. For other captures, if the target biome were monitoring of the ocean, meltwater from the mountains, or the forest, both captures would show a high likelihood of being a relevant image even though they also contain a notable cloud cover percentage.

Fig. 8 shows two images of slightly different locations within the Caspian Sea. In this image, we can again see some classification errors, most notably that *Caspiansea1* being labeled to contain *dark forest* in the upper-middle part of the image, and *Caspiansea2* being labeled to have *sand* to the right of the peninsula. Additionally, in *Caspiansea2*, some of the sand was overexposed, thus leading to being misclassified as *thick clouds*. Since the main

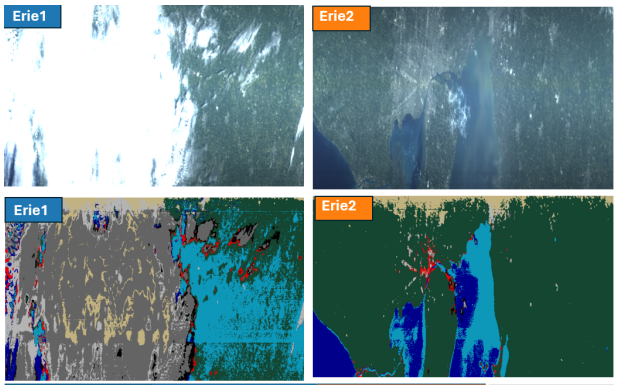


| Class | Erie1 [%] | Erie2 [%] |
|---|---|---|
| Water | 2.11 | 9.26 |
| Strange Water | 18.93 | 8.59 |
| Dark Forest | 18.76 | 76.21 |
| Urban | 1.15 | 0.66 |
| Sand | 6.43 | 4.13 |
| Thick Clouds | 9.33 | 0.74 |
| Thin Clouds | 41.1 | 0.19 |
| Shadow | 2.2 | 0.23 |

Fig. 6. Comparison of two labeled images of Lake Erie - represented by the RGB composite of the hyperspectral cube and the decoded labels - and the percentage of pixels within the different classes. Note that Erie1 has been incorrectly labeled as it shows a high amount of *strange water* that is not seen in the RGB composite of the image. Legend for the decoded labeled image is shown on the bottom right, acquired from [13].

objective of the HYPSO mission is to monitor the water quality, and both captures show a large percentage of *strange water* both would be considered good candidates for downlink.

## 6. Future Outlook

Section 5 presented results from images of the same scenes to compare pictures within the same area. The classification model proved to detect differences between captures from different times. The classification results are somewhat inaccurate, but the tests show that the concept can work with a good classification model. Going forward, two main areas for improvement are necessary: the accuracy of the classification method and class thresholds for the imaging targets.

It is necessary to retrain the classifier to be more

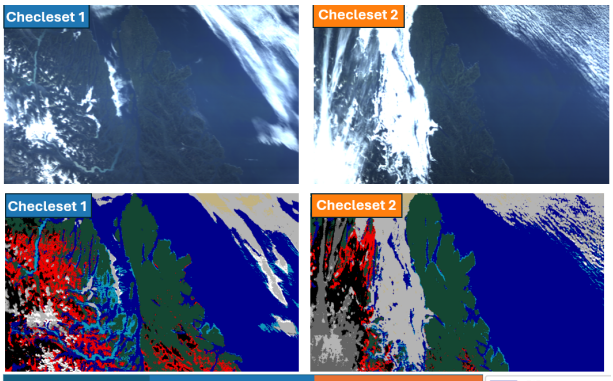


| Class | Checleset1 [%] | Checleset2 [%] |
|---|---|---|
| Water | 41.07 | 42.12 |
| Strange Water | 4.46 | 1.59 |
| Dark Forest | 25.01 | 18.8 |
| Urban | 9.25 | 2.74 |
| Ice | 1.68 | 0.25 |
| Sand | 1.22 | 1.62 |
| Thick Clouds | 9.63 | 19.56 |
| Thin Clouds | 0.95 | 6.69 |
| Shadow | 6.54 | 6.63 |

Fig. 7. Comparison of two images captured of Checleset - represented by the RGB composite of the hyperspectral cube and the decoded labels - and the percentage of pixels within the different classes. Legend for the decoded labeled image is shown on the bottom right, acquired from [13].

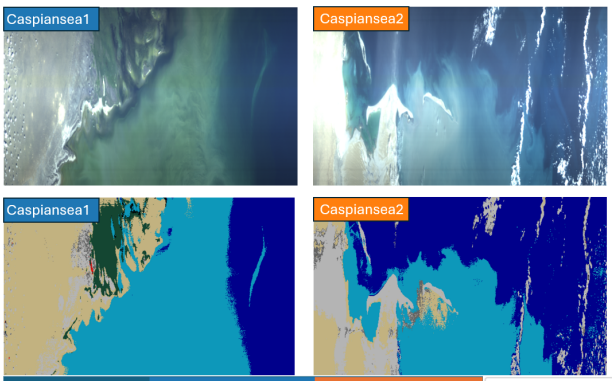


| Class | Caspiansea1 [%] | Caspiansea2 [%] |
|---|---|---|
| Water | 20.38 | 46.76 |
| Strange Water | 41.01 | 28.74 |
| Dark Forest | 4.99 | 0 |
| Sand | 28.87 | 8.98 |
| Thick Clouds | 3.92 | 13.84 |
| Thin Clouds | 0.57 | 1.59 |

Fig. 8. Comparison of two images captured of different locations within the Caspian sea - represented by the RGB composite of the hyperspectral cube and decoded labels - and the percentage of pixels within the different classes. Legend for the decoded labeled image is shown on the bottom right, acquired from [13].

accurate and change the classes to correspond better with the target biomes of imaging locations. Making new, niche classes such as *volcano* or *wildfire* is possible. Such classes could be useful in specific imaging scenarios, but adding a large number of classes can make the output more unreliable because it is more complex to differentiate between the distinct features of the classes and the potential overlap between similar classes. Having sufficient training data could also become problematic. Instead, niche imaging targets could be a better use-case for target detection methods for hyperspectral images, such as the ones described in [18]. For the HYPSO mission, having classes to separate *strange water* into more specific subclasses, such as Chlorophyll-A, could be the most beneficial, and mapping coastlines would be desirable. The classification performance could become more accurate by implementing onboard atmospheric correction and retraining the model on the corrected datasets, but that remains as future work.

Each imaging target needs certain thresholds to be considered good enough for downlink. These thresholds need to be defined after the classification model is improved due to it being too inaccurate to reliably trust the current version of it.

Once these upgrades are in place, the HYPSO-1 satellite - or other satellites with the same limitations and specifications - could be tasked with imaging a wide variety of waterbodies each day but only downlinking the ones that show any presence of algal blooms, thus increasing the satellite's coverage.

Looking further ahead, the satellite still needs to run onboard georeferencing to detect if it successfully imaged the exact targeted region. How to implement it and what tools to use requires more research and remains as future work.

## 7. Conclusion

We have presented the problem of CubeSats' limited coverage through the specific, quantified limiting factors

of the HYPSO-1 satellite. We showed our approach to enhance operations efficiency by utilising onboard classification to detect whether or not an acquired image captured the intended target biome. We presented the architectural software changes necessary to the HYPSO-1 payload, and planning and scheduling software to implement this capability, and performed testing by running onboard classification and denoting the percentage of pixels per class. The approach is deemed feasible, but requires further work in improving the classification model to make it more accurate, and in defining suitable thresholds for each imaging target for when the capture should be downlinked. With these changes, the desired outcome for the satellite is to perform more daily image acquisitions than it currently is possible to downlink. Furthermore, one can let the payload itself decide which acquisitions to downlink based on the classification output and user-set thresholds. Thus, we can increase its coverage without downlinking more captures.

**Acknowledgements**

The research leading to this work has received funding from the Research Council of Norway through from Observational Pyramid with Hyperspectral Nano-Satellites for Ocean Science (Grant No. 325961), A system-of-systems approach to real-time integrated ocean environmental monitoring (Grant No. 337495), and Mission Mjøsa. Mission Mjøsa is funded by NTNU, the county municipality of Innlandet, and the municipality of Eidsvoll, Østre Toten, Gjøvik, Lillehammer, Ringsaker, Hamar, Løten, and Stange.

We want to thank Jonas G. Røysland for his previous work and tremendous effort during his master's thesis in making the Support Vector Machine model used for onboard classification in this work.

**Contributions**

The development of the model, software, and training of the SVM was performed by Jonas G. Røysland in his previous work. The utilisation of the SVM to improve satellite operations as described in this paper, development of the changes to denote class percentages, deployment in-orbit, performing satellite operations and comparing results on the ground was performed by S.B., D.D.L., and C.C. The development of the model, software and training of the CNN was performed by J.A.J in previous work. The manuscript was reviewed by D.D.L., C.C., R.B., J.A.J, and T.A.J.

**References**

[1] A. Poghosyan and A. Golkar, "Cubesat evolution: Analyzing cubesat capabilities for conducting science missions," *Progress in Aerospace Sciences*, vol. 88, pp. 59–83, 2017, ISSN: 0376-0421. DOI: https://doi.org/10.1016/j.paerosci.2016.11.002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0376042116300951.

[2] G. Curzi, D. Modenini, and P. Tortora, "Large constellations of small satellites: A survey of near future challenges and missions," *Aerospace*, vol. 7, no. 9, 2020, ISSN: 2226-4310. DOI: 10.3390/aerospace7090133. [Online]. Available: https://www.mdpi.com/2226-4310/7/9/133.

[3] R. Birkeland, S. Berg, M. Orlandic, and J. L. Garrett, "On-board characterization of hyperspectral image exposure and cloud coverage by compression ratio," in *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2022, pp. 1–5. DOI: 10.1109/WHISPERS56178.2022.9955117.

[4] S. Berg, S. Bakken, R. Birkeland, C. Chiatante, J. L. Garrett, and T. A. Johansen, "Ground systems software for automatic operation of the HYPSO-2 hyperspectral imaging satellite," in *Sensors, Systems, and Next-Generation Satellites XXVII*, S. R. Babu, A. Hélière, and T. Kimura, Eds., International Society for Optics and Photonics, vol. 12729, SPIE, 2023, p. 1 272 912. DOI: 10.1117/12.2684263. [Online]. Available: https://doi.org/10.1117/12.2684263.

[5] N. M. Institute. "Met weather api." Visited 2024-09-18. (), [Online]. Available: https://api.met.no/.

[6] G. Giuffrida *et al.*, "The φ-sat-1 mission: The first on-board deep neural network demonstrator for satellite earth observation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2022. DOI: 10.1109/TGRS.2021.3125567.

[7] ESA. "Next artificial intelligence mission selected." (), [Online]. Available: https://www.esa.int/Applications/Observing_the_Earth/Ph-sat/Next_artificial_intelligence_mission_selected.

[8] ESA. "Chime onboard processing: Cloud detection and selective compression." (), [Online]. Available: https://hyperspectral2022.esa.int/iframe-agenda/files/presentation-174.pdf.

[9] A. V. et al. "Platform-agnostic hyperspectral model for onboard cloud detection." presentation at the 2023 CubeSat Winter Workshop in Aalto, Finland. (Jan. 2023), [Online]. Available: https://dragon3.esa.int/eogateway/documents/d/earth-online/05_vh-roda_2023_vandenhoeke.

[10] R. Birkeland, S. Berg, M. Orlandic, and J. L. Garrett, "On-board characterization of hyperspectral image exposure and cloud coverage by compression ratio," in *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2022, pp. 1–5. DOI: 10.1109/WHISPERS56178.2022.9955117.

[11] J. A. Justo, J. L. Garrett, M.-I. Georgescu, J. Gonzalez-Llorente, R. T. Ionescu, and T. A. Johansen, *Sea-land-cloud segmentation in satellite hyperspectral imagery by deep learning*, 2023. arXiv: 2310.16210 [cs.CV]. [Online]. Available: https://arxiv.org/abs/2310.16210.

[12] J. G. Røysland, D. D. Langer, S. Berg, M. Orlandić, and J. L. Garrett, "Hyperspectral classification onboard the hypso-1 cubesat," in *2023 13th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2023, pp. 1–5. DOI: 10.1109/WHISPERS61460.2023.10431215.

[13] J. G. Røysland, "Real-time classification onboard the hypso-1 satellite," M.S. thesis, Norwegian University of Science and Technology, https://hdl.handle.net/11250/3091477, 2023.

[14] D. D. Langer, T. A. Johansen, and A. J. Sørensen, "Consistent along track sharpness in a push-broom imaging system," in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023, pp. 4486–4489. DOI: 10.1109/IGARSS52108.2023.10283310.

[15] S. Bakken *et al.*, "Hypso-1 cubesat: First images and in-orbit characterization," *Remote Sensing*, vol. 15, no. 3, 2023, ISSN: 2072-4292. DOI: 10.3390/rs15030755. [Online]. Available: https://www.mdpi.com/2072-4292/15/3/755.

[16] C. Chiatante, D. D. Langer, J. L. Garrett, R. Birkeland, S. Berg, and M. Orlandić, "Onboard hyperspectral classification enables georeferencing," in *2023 13th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2023, pp. 1–5. DOI: 10.1109/WHISPERS61460.2023.10430746.

[17] G. Rabideau *et al.*, "Mission operations of earth observing-1 with onboard autonomy," in *2nd IEEE International Conference on Space Mission Challenges for Information Technology (SMC-IT'06)*, 2006, 7 pp.–373. DOI: 10.1109/SMC-IT.2006.48.

[18] Đ. Bošković, M. Orlandić, and T. A. Johansen, "A reconfigurable multi-mode implementation of hyperspectral target detection algorithms," *Microprocessors and Microsystems*, vol. 78, p. 103 258, 2020, ISSN: 0141-9331. DOI: https://doi.org/10.1016/j.micpro.2020.103258. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S014193312030418X.