ON-BOARD CHARACTERIZATION OF HYPERSPECTRAL IMAGE EXPOSURE AND CLOUD COVERAGE BY COMPRESSION RATIO

R. Birkeland, S. Berg and M. Orlandic

Norwegian University of
Science and Technology
Department of Electronic Systems
O. S. Bragstad plass 2A,
Trondheim, N-7034, Norway

J. L. Garrett

Norwegian University of
Science and Technology
Department of Engineering Cybernetics
O. S. Bragstad plass 2D,
Trondheim, N-7034, Norway

ABSTRACT

Hyperspectral images are useful for remote sensing applications due to the fine spectral resolution relative to regular RGB- and multi-spectral images. The images can be used to study *what* is observed, rather than just *that* a given feature is observed. The high spectral resolution comes with a challenge for both small and big satellites: the data volume per observation. For a small earth observation satellite, the time and energy to downlink data is one of the main driving factors behind data latency and imaging capacity.

Effective use of satellite time calls for a smart imaging pipeline, both when it comes to planning, execution, and processing of data from observations. The imaging and processing pipeline needs to utilize both on-board and on-ground processing tools and steps. In this paper, a set of the first observations from the HYPSO-1 satellite is analyzed to evaluate how quickly and effectively the compression ratio can be used to identify observations that should be prioritized for downlinking, and which data sets seem to be of less value due to over/under exposure or cloud coverage.

***Index Terms*—** HYPSO-1, CubeSat, CCSDS-123, Operations

1. INTRODUCTION

Hyperspectral cameras are an invaluable tool for resource and environmental monitoring since they can reveal the physical properties of the observed targets in addition to their spatial features. The hyperspectral data is used to study the land cover composition at a level of spectral detail. Example systems and instruments are HICO, HyperScout, PRISMA, EnMAP, PACE. Small satellites, such as CubeSats, are unique platforms for carrying such cameras for countries and institutions that wish to monitor their surroundings but do not have the capital to invest in larger satellites. The well-established new-space industry can offer mission-tailored and cost-effective satellite platforms. Equipped with suitable imaging payloads, the satellites can be pointed to locations of interest. In addition, mission-specific satellites can provide more frequent target passes and more agile mission planning, such as choosing/prioritizing targets depending on weather conditions, end user demands or guiding from other sensor systems.

CubeSats, which are small satellites with a standardized form factor [1], have revolutionized the access to space for new actors as universities, research organizations, startups, and even private groups are building and launching their own satellites. Many CubeSat mission groups use an agile approach to adapt operations during a mission. Planning for a software architecture that enables over-the-air updates allows for adding new functionality to the mission throughout its lifetime [2], so it is possible to adjust the payload behavior during the mission. More flexible operations allow operators to “order” observations on-the-fly with a low lead-time (down to less than an hour), depending on the time of day and the weather conditions. Several hyperspectral CubeSats have been launched to date, like HYPerspectral Smallsat for Ocean observation (HYPSO)-1 [3], Aalto-1 and Hyperscout [4, 5].

Because of the detailed spectral information hyperspectral images contain, they are large and lead to significant demands on the downlink capacity of many small satellites, even when compressed. Therefore, there are choices about how images should be prioritized in the downlink queue and which images could be discarded. Selection and prioritization of images to downlink are important when there are many areas-of-interest to monitor compared to the storage available on the satellite.

In this proceeding, we investigate whether the compression ratio of the hyperspectral images can be used as a metric for assessing the *quality* of an image. We discuss the first in-flight experiences of the small HYPSO-1 satellite, particu-

---

This work was supported by the Research Council of Norway, Equinor, DNV GL and Sintef through the Centers of Excellence funding scheme, Grant 223254 - Center for Autonomous Marine Operations and Systems (AMOS) and the Research Council of Norway through the IKTPLUSS programme grant 270959 (MASSIVE).

larly how its operations can be made more autonomous during its operational phase. We compare datasets collected by the HYPSO-1 CubeSat's hyperspectral imager over the course of the two first months in orbit (early February to early April 2022), the payload commissioning phase.

2. BACKGROUND

2.1. Hyperspectral image quality

The ocean surface is dark compared to nearby land and clouds. Thus, when the camera settings are tuned to capture ocean color with a high signal-to-noise ratio, many pixels that contain clouds become saturated unless the pixels have a very high dynamic range. When one band of a pixel becomes saturated, the spectral information in that pixel becomes unusable.

Saturated pixels constitute the main limit on image quality assessed in this proceeding. Similarly, if the satellite points away from the earth, many pixels will collect no photons, a similar detriment to image quality. During the study, additional assessments of image quality, such as virtual dimensionality, were also tested [6]. However, at the level of operations, the virtual dimensionality did not provide additional insight into whether an or not an image should be downlinked.

2.2. Image Compression

Consultative Committee for Space Data System 123 (CCSDS123) is a lossless compression algorithm for multi- and hyperspectral imagery [7] that relies on a simple, predictive encoding. On HYPSO-1, the hyperspectral processing pipeline includes both software and FPGA implementations [8] of the CCSDS-123 Issue 1 algorithm. The FPGA IP core supports the Band-Interleaving-by-Pixel (BIP) ordering and the majority of user-defined compression parameters defined by the standard. In HSI remote sensing, it is common for no loss of information in the experiments to be accepted, so that lossless compression is preferable. The CCSDS-123 Issue 1 standard [9, 7], targeting multispectral and hyperspectral images, is characterized by low complexity and a high compression ratio that reduces demand for down-link time and thus also saves onboard energy. In this way, the satellite can perform more observations.

Saturated pixels affect the compression ratio when they are in close spatial or spectral proximity. In that case, the predictive compression algorithm does not need to include as much information because the pixels are identical. While a similar low compression ration could be expected by relatively uniform scenes shot noise does not appear in saturated pixels and thus a lower compression ratio is expected for them even relative to the open ocean.

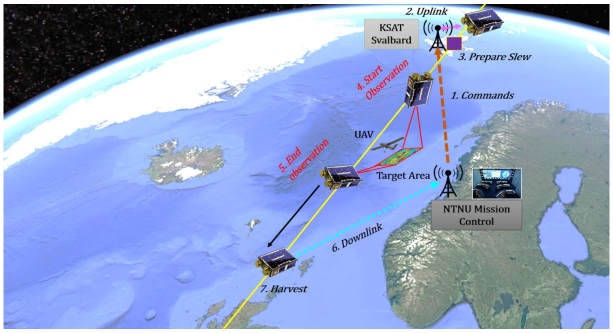

Fig. 1. HYPSO-1 operational concept.

2.3. HYPSO-1 mission

The HYPSO-1 satellite is the first scientific CubeSat developed at a Norwegian university [3]. The satellite bus is procured commercially from the Lithuanian company NanoAvionics, the payload consisting of an push-broom imager and the on-board-processing system is designed and made by Norwegian University of Science and Technology (NTNU) researchers and students [10]. HYPSO-1 records 956 lines in one observation for the normal imaging mode. Exposure and frame rate can be adjusted. Early tests were performed with fewer lines.

The operational concept is shown in figure Fig. 1. Imaging parameters such as start time, pointing, exposure, and duration are uploaded prior to a pass over an area of interest. Following the observation, the image is downloaded to the mission control center.

Fig. 2. The sizes of the data captured by the HYPSO-1 satellite during the first two months of imaging.

2.4. Operational constraints

CubeSats are powerful assets despite their small size, but with small size there are several constraints which impact its operations. The two most important constraints are *available en-*

ergy and *downlink capacity*, especially for Earth Observation (EO)-missions which generate a lot of data.

Academic projects often encounter other resource constraints as well, including limited funding and human resources. Because operational tasks are often distributed amongst team members, a high degree of automation is needed, as well as confidence in the automated processes, to ensure the best utilization of a satellite. Thus, academic projects in particular stand to benefit from automated dataflows and pre-processing of observations.

2.4.1. Downlink capacity

For HYPSO-1, the nominal imaging mode produces an hyperspectral imager (HSI)-cube of 153 MB. Even with a high-rate S-band downlink of more than 1 Mbps datarate, downlinking a full cube will take more than 20 minutes. The radio link is shared with satellite operations such as telemetry downlinking and command uploading. The NTNU ground station sees the satellite for an average of ten passes a day, with a total theoretical duration of 100 minutes. Of this, approximately 60 minutes could be available for downlinking hyperspectral data if the operation is streamlined and error-free. This corresponds to ≈450 MB of data per day, which in turn corresponds to three uncompressed HSI cubes. When a target location is near the ground station, about 10 minutes of possible downlink time is lost to imaging.

In HYPSO-1, lossless CCSDS123 compression reduces the size of the data to be downlinked to 80 MB or less, depending on the contents of the image cube. Still, it will require two or even three ground station passes to complete the downlink of one cube. Operationally, it is therefore important to ensure that the data downlinked is valid and valuable. Compared to downlinking uncompressed cubes, the ideal operation of the ground station can then allow for downlinking between four to six compressed HSI cubes.

2.4.2. Planning and scheduling

HYPSO-1 is commanded to make observations over selected target areas. An imaging schedule must be uploaded to the satellite before the time that the satellite will pass a target. In practice, the delay between the upload and the imaging can be as little as a few minutes (if a ground station is near the target area) but can also be set up several days in advance.

During the first period of operation of HYPSO-1, the weather forecasts were not used deciding the imaging schedule for most targets because the scheduling involved a great deal of manual work prior to the development of better planning tools. Later, daily imaging over a selected target was set up, also ignoring the potential cloud coverage. In addition, imaging was set up over targets-of-opportunity every day to collect a variety of datasets, also ignoring any cloud coverage. Due to this, several of the captured scenes contained a high number of saturated pixels.

As operational procedures become more refined, two things are expected to change. First, when the manual work of selecting a target and defining the image schedule is less demanding, we expect to reduce the minimum delay between making an image request and the imaging itself. Once that delay is reduced, assessments of the weather forecast, cloud coverage, in particular, will become an essential parameter for imaging planning. With image planning immediately prior to the imaging, the confidence in weather predictions is higher, so there will be a lower risk of cloud coverage ruining the image quality compared to the first weeks of operations when weather forecasts was given less attention.

Second, for agile and "short-timed" operations combined with the desire to make as many observations per day, it will be increasingly important to quickly assess if a data-set is worth downlinking or not. One possible parameter to assess is the compression ratio of the HSI-cube may give a strong indication of whether an observation should be downlinked or not.

2.4.3. Energy budget considerations

All satellites have limited energy storage onboard and a finite capacity of energy harvesting. The most energy-intensive operations for HYPSO-1 are connected to data management. While the imaging operation itself is a short duration operation, preparing the data for downlink and downlinking it are more energy demanding tasks. Ultimately, satellites such as HYPSO-1 are energy limited by the mentioned factors, thus driving the possible imaging duration and the number of observations per day. It is possible to extend the downlink time by adding more ground stations to the system. However, at some point, the energy budget would become negative. Operating HYPSO-1 with one ground station pass per orbit keeps the power budget positive with a margin.

3. ANALYSIS AND RESULTS

In this section, we compare a set of the 59 first captures performed by HYPSO-1 during the first weeks of operation. The images are captured over different areas of the world, at different local times, and some with different exposure settings.

3.1. Method

For all datasets, which may contain a varying number of exposed frames we calculated the *compression ratio*, r, whic is defined as the compressed file size divided by the original size of the image.

A low-res RGB composite was generated, and a visual assessment of those RGB images was carried out to classify them into *usable* or *not usable* images, based on the amount of saturated/white pixels in the scene. The visual assessment was then linked to the calculated compression ratio to investigate if there is a threshold to determine the quality of a hyperspectral image. In addition, the portion of saturated pixels in each scene was calculated, and compared with the compression ration.

3.2. Results

Figure 3 shows an example of a scene with primarily saturated pixels, with a low r, whereas Figure 4 shows a scene over the ocean with only a few clouds and with more visible ocean.

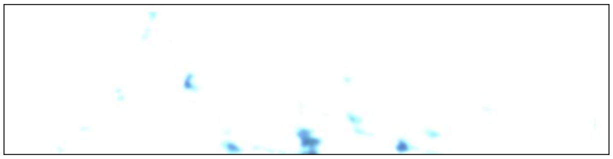

Fig. 3. Low resolution RGB-render of scene with mostly white pixels, $r = 0.289$. A border has been added to the image for clarity. Capture date 2022-02-06.

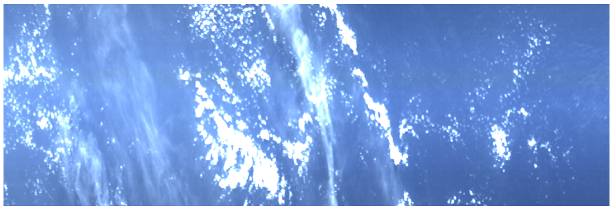

Fig. 4. Low resolution RGB-render of scene of ocean with scattered clouds. $r = 0.498$. Capture date 2022-03-30.

The compression ratio and saturated pixels for all images are compared and plotted in Figure 5.

There are other possible causes for low compression ratios in addition to cloud coverage. Sometimes images can also be negatively affected by pointing errors. Figure 6 shows an image where the pointing did not align with the target area. This also reduces the value of an observation.

4. DISCUSSION

We observe that compression rations lower than 0.35 are not worthwhile for in-depth analysis from the assessed dataset. Ratios up to 0.42 may contain a lot of clouds, but small portions of the image may show the ocean. Ratios over 0.45 contain more diverse information. Therefore, moving forward, we are planning to automatically discard data-sets with a compression ratio of less than 0.35. Note, however, that above 0.4 the compression ratio does not seem to correlate strongly with image quality.

Beyond satellites, these results could assist in other situations where there is limited connectivity between the imaging platform and the operator. For example, hyperspectral imaging from high-altitude drones and underwater vehicles. The advantages of looking at a compression ratio rather than directly counting saturated pixels include the re-use of code and limiting different tasks to be completed on board. The results presented here pertain to the CCSDS-123 algorithm because it had been developed for processing on-board HYPSO-1. However, it could also be worth studying different compression algorithms as well, perhaps including lossy compression.

There are two ways of incorporating this information into in the processing pipeline. Either the metadata about a capture can be downlinked for analysis on ground, or the satellite can make its own decision utilizing the on-board processing capabilities. If an operator analyzes metadata before deciding whether to downlink the data, the latency will increase. There will be one pass for downlinking the metadata, then time for analyzing it, before uploading the decision (delete or keep) on the following pass. The extra pass adds at least 1.5 hours of latency (time between two passes). Then, the satellite must prepare the data for downlinking before the data may be downlinked on the next pass, at least three hours prior to the observation. On the other hand, if the satellite itself decides

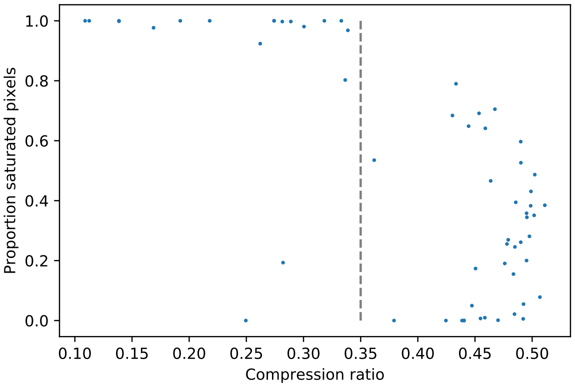

Fig. 5. The proportion of pixels which are saturated, as a function of compression ratio. The two images with few saturated pixels and a low compression ratio either had a very low signal or pointed in the wrong direction, imaging space rather than the earth. The dashed line indicates the proposed downlink limit.

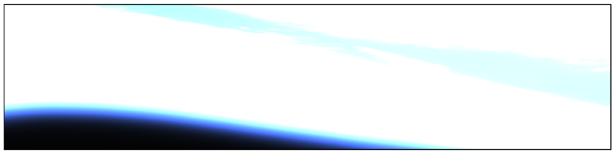

Fig. 6. Low resolution RGB-render of a capture meant to be off the coast of Norway, but suffered from a pointing error. $r = 0.303$. Capture date 2022-04-25.

to keep the observation, it can be prepared for downlinking immediately and downlinked on the next possible pass.

As shown in this paper, the compression ratio is a parameter that gives insight into the quality of the capture, but it does not address *why* the quality is of that level. Most likely, a low compression rate is caused by a cloud cover, but as Figure 6 shows, it can also be caused by pointing errors. Thus, relying entirely on a measure that only measures the quality and not *why* the quality is as it is might hide underlying issues.

A more extensive alternative will be to let the satellite downlink an RGB composite of the hyperspectral observation if the first assessment is *discard*. Then an operator can decide to downlink the observation depending on a visual assessment, as shown for Figures 3 and 4. To ensure that the operator will still be informed about possible underlying issues, the RGB composite could always be downloaded, while the full hyperspectral observation is discarded. This would use the compression rate parameters strength, which is to determine the quality of the capture, while handling its weakness by letting a human determine *why* the capture was of bad quality.

5. CONCLUSIONS

In this paper we show that using the simple metric of the compression ratio of a hyperspectral image as an indicator of whether an image should be downlinked or not can save time and energy for the satellite. As much as 20 minutes of downlink time and about 3.6 Wh energy could be saved per occasion. Letting the satellite decide whether or not to downlink an observation could also reduce latency, typically 2 to 3 hours for the good images.

Acknowlegdements

The authors wish to express their thanks to the whole HYPSO team for making the satellite possible, particularly the software group, represented by Sivert Bakken and Dennis Langer, and the overarching vision provided by Tor Arne Johansen and Evelyn Honoré-Livermore.

6. REFERENCES

[1] Hank Heidt, Jordi Puig-Suari, Augustus S. Moore, Shinichi Nakasuka, and Robert J. Twiggs, “CubeSat: A New Generation of picosatellite for education and industry low-cost space experimentation,” in *Proc. 14th Annual AIAA/USU Conference on Small Satellites*, 2000, Lessons Learned - In Success and Failure, SSC00-V-5, pp. 1–19, `https://digitalcommons.usu.edu/smallsat/2000/All2000/32/`.

[2] Sivert Bakken, Evelyn Honore-Livermore, Roger Birkeland, Milica Orlandic, Elizabeth F. Prentice, Joseph L. Garrett, Dennis D. Langer, Cecilia Haskins, and Tor A. Johansen, “Software Development and Integration of a Hyperspectral Imaging Payload for HYPSO-1,” in *2022 IEEE/SICE International Symposium on System Integration (SII)*, 2022.

[3] Mariusz E. Grøtte, Roger Birkeland, Evelyn Honoré-Livermore, Sivert Bakken, Joseph L. Garrett, Elizabeth F. Prentice, Fred Sigernes, Milica Orlandić, J. Tommy Gravdahl, and Tor A. Johansen, “Ocean color hyperspectral remote sensing with high resolution and low latency–the hypso-1 cubesat mission,” *IEEE Transactions on Geoscience and Remote Sensing*, pp. 1–19, 2021.

[4] Marco Esposito and A Zuccaro Marchi, “In-orbit demonstration of the first hyperspectral imager for nanosatellites,” in *International Conference on Space Optics—ICSO 2018*. International Society for Optics and Photonics, 2019, vol. 11180, p. 1118020.

[5] M. Rizwan Mughal, J. Praks, R. Vainio, P. Janhunen, J. Envall, A. Näsilä, P. Oleynik, P. Niemelä, S. Nyman, A. Slavinskis, J. Gieseler, N. Jovanovic, B. Riwanto, P. Toivanen, H. Leppinen, T. Tikka, A. Punkkinen, R. Punkkinen, H.-P. Hedman, J.-O. Lill, and J.M.K. Slotte, “Aalto-1, multi-payload cubesat: In-orbit results and lessons learned,” *Acta Astronautica*, vol. 187, pp. 557–568, 2021.

[6] Chein-I Chang and Qian Du, “Estimation of number of spectrally distinct signal sources in hyperspectral imagery,” *IEEE Transactions on geoscience and remote sensing*, vol. 42, no. 3, pp. 608–619, 2004.

[7] Consultative Committee for Space Data Systems, *Lossless Multispectral and Hyperspectral Image Compression-CCSDS 123.0-B-1*, CCSDS Secretariat, Washington, DC, USA, 2012.

[8] Milica Orlandić, Johan Fjeldtvedt, and Tor Arne Johansen, “A parallel fpga implementation of the ccsds-123 compression algorithm,” *Remote Sensing*, vol. 11, no. 6, 2019.

[9] Consultative Committee for Space Data Systems, *Lossless Multispectral and Hyperspectral Image Compression-CCSDS 120.2-G-1*, CCSDS Secretariat, Washington, DC, USA, 2015.

[10] Elizabeth Frances Prentice, Mariusz Eivind Grøtte, Fred Sigernes, and Tor Arne Johansen, “Design of a hyperspectral imager using COTS optics for small satellite applications,” in *International Conference on Space Optics — ICSO 2020*, Bruno Cugny, Zoran Sodnik, and Nikos Karafolas, Eds. International Society for Optics and Photonics, 2021, vol. 11852, pp. 2154 – 2171, SPIE.