Tackling the Satellite Downlink Bottleneck with Federated Onboard Learning of Image Compression

Pablo Gómez[°†‡] pablo.gomez@esa.int

Gabriele Meoni[*†⋆‡] gabriele.meoni@esa.int

[*]Φ-lab, ESA/ESRIN [⋆] Department of Space Engineering, Aerospace Engineering, TU Delft [†] AI Sweden, Göteborg, Sweden
[°]Advanced Concepts Team, ESA/ESTEC [‡] Shared first authorship

Abstract

Satellite data transmission is a crucial bottleneck for Earth observation applications. To overcome this problem, we propose a novel solution that trains a neural network on board multiple satellites to compress raw data and only send down heavily compressed previews of the images while retaining the possibility of sending down selected losslessly compressed data. The neural network learns to encode and decode the data in an unsupervised fashion using distributed machine learning. By simulating and optimizing the learning process under realistic constraints such as thermal, power and communication limitations, we demonstrate the feasibility and effectiveness of our approach. For this, we model a constellation of three satellites in a Sun-synchronous orbit. We use real raw, multispectral data from Sentinel-2 and demonstrate the feasibility on space-proven hardware for the training. Our compression method outperforms JPEG compression on different image metrics, achieving better compression ratios and image quality. We report key performance indicators of our method, such as image quality, compression ratio and benchmark training time on a Unibap iX10-100 processor. Our method has the potential to significantly increase the amount of satellite data collected that would typically be discarded (e.g., over oceans) and can potentially be extended to other applications even outside Earth observation. All code and data of the method are available online to enable rapid application of this approach.

1. Introduction

Earth observation data are becoming increasingly valuable for a variety of applications such as environmental monitoring, disaster management, urban planning and security [24, 28, 33]. These applications are especially relevant for addressing the challenges posed by climate change. The number of Earth observation satellites is also growing rapidly, as is the amount of data produced. For instance, Copernicus satellites alone produce around 20 TB of data daily[1]. However, transmitting large amounts of data from satellites to ground stations is an increasingly challenging task due to the limited availability of ground station windows and the limited bandwidth of communication links [4, 29]. Therefore, there is a need for efficient methods to compress satellite data and reduce the transmission load.

Numerous approaches have been proposed to tackle this problem. Some research works leverage Deep Neural Networks (DNNs) to filter data to transmit to the ground by limiting the transmitted data to actionable information only [10, 15, 16, 32]. Examples of actionable information include the coordinates of events of interest or cloud-free images. This approach is particularly suitable for latency-constrained applications and early-alert systems [10, 11] but does not translate well to scientific missions aiming to provide the end users with as much data as possible.

For those missions, alternative approaches relying on lossless or near-lossless compression schemes are preferred [22]. In this way, payload data are compressed on board the satellite before downlink and decompressed at the ground segment with no or limited data loss. In particular, solutions leveraging autoencoders have been proposed [2, 18] to learn to encode the input data into a lower-dimensional representation and decode them back to the original data. Although autoencoders have been shown to outperform JPEG and other onboard payload compression methods in terms of compression factors/quality trade-offs, previous approaches assumed training a model on ground, which is often infeasible as it assumes the data are already present and is not targeted at raw data available on a satellite.

This paper proposes an initial investigation for a novel mission concept "LICOS" (Learning Image Compression

[1]New interface makes open Earth Observation data truly open. Available online at: https://projects.research-and-innovation.ec.europa.eu/en/projects/success-stories/all/new-interface-makes-open-earth-observation-data-truly-open. Accessed: 2024-04-12

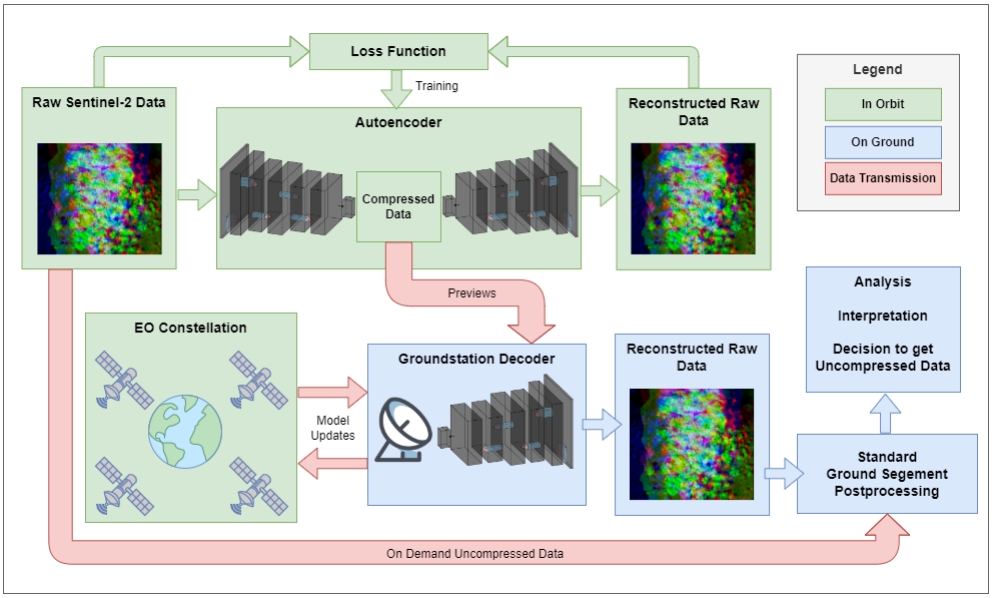

Figure 1. Schematic overview of the proposed method; Notice that the method enables sending down heavily compressed previews of images with minimal communication overhead.

On board a Satellite constellation) that capitalises on an Artificial Intelligence (AI)-based compression scheme to enable data selection by the end-user based on compressed previews. To this aim, we leverage satellites' increasing onboard computing capabilities to train an autoencoder that can compress raw data and only send down heavily compressed previews of the images. The compressed previews can then be used to select only specific images that are interesting for further analysis and request the larger, raw images on demand. This way, the transmission can be prioritized according to the user's needs and preferences. Thus, this method can provide huge data transmission savings and overcomes the risks of losing relevant data typical of standard onboard processing techniques [1, 15, 16]. A schematic of the approach is given in Fig. 1.

However, training an autoencoder on board a satellite constellation involves several operational constraints such as thermal, power, bandwidth and communication limitations that affect the performance and feasibility of the learning process [17]. Moreover, if multiple satellites are involved in the same mission, they shall coordinate and synchronize their learning activities in a distributed fashion using federated or decentralized learning techniques [25, 30]. These aspects are often neglected or oversimplified in existing works on satellite image compression.

To model the impact of thermal, power and communication constraints and optimize the learning process under realistic conditions, we use PASEOS. This open-source Python module simulates the environment to operate multiple spacecraft [17]. We also exploit multiple satellites' computing and data collection capabilities and train the autoencoder in a distributed way using federated learning techniques [17, 25, 30]. These techniques are ideal for our problem since they do not require labels or centralized coordination for the data. In contrast to existing works on satellite image compression [1, 2, 18], we account for various operational constraints in our approach and demonstrate its feasibility and effectiveness directly on raw data available on board, taking a significant step to close the reality gap of many machine learning approaches. We demonstrate our approach on a hypothetical constellation of three 6U CubeSats equipped with Unibap iX-10 100 satellite processors [6] in Sentinel-2-like orbits, a constellation of two polar-orbiting satellites that monitor land surface conditions as part of ESA's Copernicus program [12]. We use real, multispectral raw data from Sentinel-2 that would be available on a satellite (in contrast to commonly used, post-processed data such as the EuroSAT dataset [20, 21]) and compare our compression method with JPEG compression on different metrics such as peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), bits per pixel (BPP).

The main contributions of this paper are:

- We present the first work demonstrating the feasibility of training an image compression neural network on board satellites.
- We consider the operational and physical constraints of the spacecraft using PASEOS and show how they affect the learning process.
- We use real raw data from Sentinel-2 and a Unibap iX-10 100 satellite processor to demonstrate the feasibility and

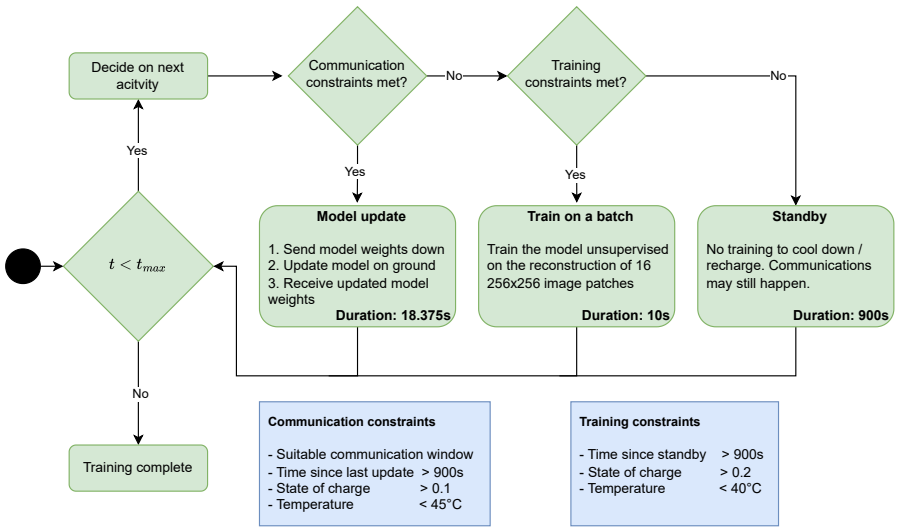

Figure 2. Flowchart describing the logic to select the next spacecraft activity to be performed. Possible activities include training, updating the model and being on standby.

effectiveness of our approach. On the Unibap iX-10 100 device, training one batch of 16 $256px \times 256px$ takes 10 s and 8 s respectively when 13 bands and one band are used.

- We compare our compression method with JPEG compression on different metrics and show that we achieve better compression ratios and image quality.

We believe that our method can transform current Earth observation paradigms by allowing the collection of significantly more data that would typically be discarded (e.g., data collected over oceans) and reducing the required data transmission and thus freeing up communication windows. Moreover, our method can conceivably also be applied to other modalities and domains beyond Earth observation. All code used here is available open-source online. [2]

2. Materials and Methods

In the remainder of the manuscript, we describe how we model the satellites and the machine learning approach to compress the raw imaging data, detail the approach to distribute training amongst the satellites and give detailed specifications of the software and hardware setup. An overview of the training algorithm is given in the flowchart in Fig. 2.

2.1. Satellite Modeling

For many machine learning applications in the space sector, practical considerations, such as the physical constraints of spacecraft in terms of thermal, power and communications budgets, are often neglected or treated superficially. However, the harsh space environment is often the decisive factor for whether something is viable. Thus, we rely on a Python module called PASEOS [17], which has been demonstrated previously to be well-suited to allow consideration of a variety of operational constraints in onboard machine learning applications [26].

With PASEOS, we model Keplerian orbits of the constellation and compute windows with ground stations based on the satellites' angle over the horizon. Communication with the ground stations is limited by a fixed available bandwidth during the window. We address the power budget by modeling the battery state-of-charge (SoC) based on consumption from training, communication, and standby usage and charging outside eclipses through solar panels. A single-node ordinary differential equation model is used for the thermal model of the satellites as provided by PASEOS [17]. Overall, a PASEOS simulation runs in parallel and asynchronously with the neural network training using the Message Passing Interface (MPI). To avoid MPI ranks advancing quicker than others, synchronization of the simulation time of each rank was performed every 10 minutes of simulation time. A detailed overview of the specific parameters of the constraints we tested is given in 3. To obtain realistic values for training times, we profiled the performance of the training on a Unibap iX10-100 processor. Communication modeling assumed a downlink data rate of 10 Mbit/s.

2.2. Neural network architecture and training algorithm

Several recent works have demonstrated great advances in image compression with neural networks [4, 5, 13, 14, 19],

[2] https://github.com/gomezzz/LICOS. Accessed: 2024-04-12

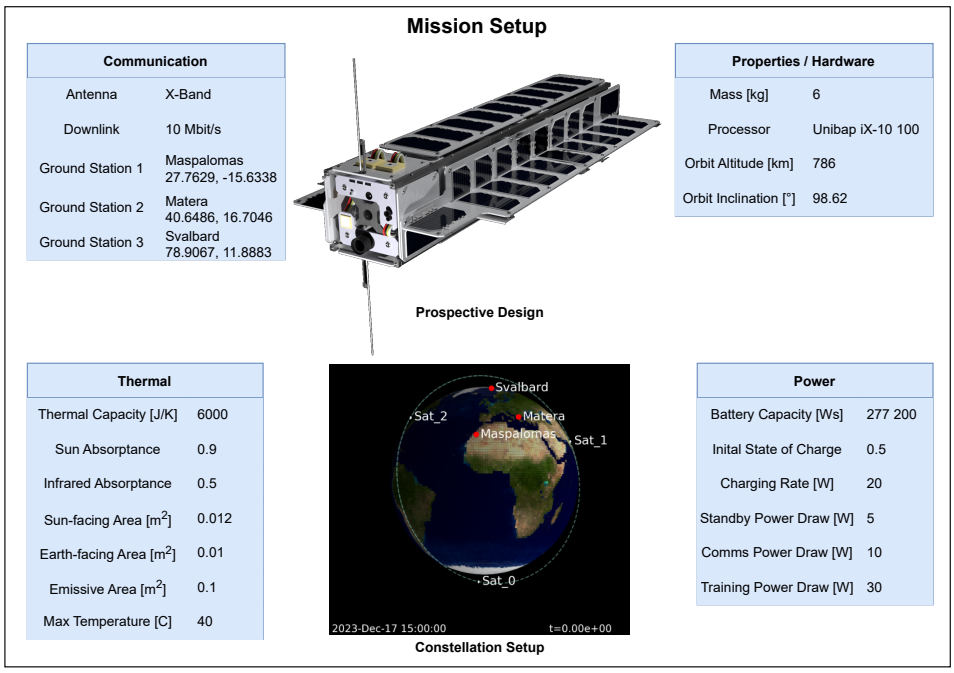

Figure 3. Mission setup and prospective satellites design.

outperforming classic compression techniques. These types of techniques have seen limited application for satellite imagery, where only a small number of works have explored so-called neural compression techniques [1, 3, 8].

In this work, we rely on the model proposed by Ballé *et al*. [4], which is implemented in the PyTorch framework CompressAI [5]. We modified the proposed model to raw, multispectral Sentinel-2 data by changing the parameters of the synthesis and analysis transforms g_s and g_a, as well as the entropy bottleneck of the model.

Pretrained models and most established compression methods focus on grayscale or RGB images. This is one of the challenges as many satellites gather multispectral data with a larger number of image channels [7, 12, 16]. The data from Sentinel-2 used in this work are 13-channel images. Consequently, we study different schemes for the compression: One called *merged*, where we compress all 13 channels jointly, and one called *split*, where the image channels are compressed individually. Note that even in split we use the same model for all 13 channels to avoid additional overhead. Conceivably, training one model for each channel may provide additional benefits.

2.3. Federated learning protocol and communication scheme

Even though our method can be applied on individual satellites, constellations provide a crucial benefit to our proposed method as they allow training in parallel on multiple devices that also collect data from different places simultaneously [23].

To facilitate this, we propose using the ground stations of the constellation as the central server in a federated learning setup. During communication windows with the ground stations, the satellites send their model to the ground station and receive an updated model from it. Note that this is quite efficient as the proposed models are smaller (11.8 MB) than even a single image (49.12 MB), introducing only minimal communications overhead. Alternative satellite network patterns relying satellite communicating with inter-satellite links are possible, as described in our previous work [26], which are not investigated in this work.

On the ground, we use a modified version of federated averaging, in which we weigh the update based on the model performance on a validation set of images. Thus, given the validation losses $\mathcal{L}_{sat}$ and $\mathcal{L}_{central}$ of the satellite and central model on the ground, respectively, the updated model weights ω are calculated as in Eq. 1:

$$\omega = \frac{\mathcal{L}_{sat}}{\mathcal{L}_{sat} + \mathcal{L}_{central}} \omega_{central} + \frac{\mathcal{L}_{central}}{\mathcal{L}_{sat} + \mathcal{L}_{central}} \omega_{sat} \quad (1)$$

where $\omega_{central}$ and ω_{sat} are the current model weights of the central server and satellite, respectively. In this fashion, we push the central model to converge towards better solutions.

The updated model is then sent back to the satellite for further training. The exchange of the models is constrained by the satellite modeling performed with PASEOS,

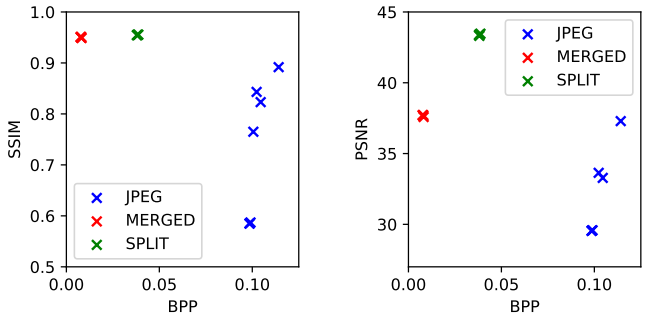

Figure 4. SSIM versus BPP and PSNR vs SSIM trends for JPEG, merged, split compression algorithms. Results for "Split" and "Merged" models, which were tested for quality factors of 1,2,4,8, overlap in the image and look like a single point.

i.e. available communication windows and bandwidth. The time for computing the validation loss on the ground is neglected. Parallel training is made feasible by allocating one training device per MPI rank.

3. Results

As a demonstration of the feasibility and advantages of the proposed approach, we showcase results from a comprehensive simulation of three satellites in low-Earth orbit jointly learning to compress raw Sentinel-2 data while balancing thermal, power and communication constraints using real satellite processors. We compare results with JPEG, an established compression standard, and explore a mono-channel (*split*) as well as multi-channel (*merged*) approach.

3.1. Setup

3.1.1 Dataset

Many works related to onboard machine learning applications for satellite imagery rely on high-end products, which underwent dedicated processing including orthorectification, geometric and radiometric correction [31]. This is quite a strong assumption as small satellites may lack the processing power for this [10]. For this reason, we instead demonstrate our method on a subset of the THRawS dataset consisting of raw Sentinel-2 imagery[9]. The dataset contains over 900 Sentinel-2 raw granules depicting volcanic eruptions, wildfires, and other areas free of thermal anomalies. A raw granule corresponds to a portion of $23 \times 25 km^2$ acquired by a single onboard detector. The "raw" format is not a standard Sentinel-2 product,[3] but corresponds to decompressed Level-0 data with additional metadata [10].

The following results assume a split of the dataset into train, validation and test sets with a ratio of 0.7, 0.1 and 0.2, respectively. A total of 860 images with resolutions ranging from $384px \times 1296px$ up to $2304px \times 2592px$ were used. Specifically, we randomly select one $256px \times 256px$ patch from each train image and pick the central $256px \times 256px$ crop for each validation image to speed up the training and limit the model memory footprint. The images in the test set are from different geographical locations than train and validation to avoid data leakages and ensure generalizability to different spatial locations.

Data were prepared using PyRawS. [4]. The latter is an open-source Python package designed to open, visualize, and post-process raw Sentinel-2 data, for instance, by applying lightweight bands coarse registration and georeferencing processes. In the frame of this project, PyRawS was only used to export the various data in a format ready to use for training, including information on granule locations needed for the train/test splitting based on geographical locations. Since the Sentinel-2 satellites have bands of different resolutions, i.e., 10 m, 20 m, 60 m, for the "merged" models we resampled all the bands to 20 m spatial resolution to ensure that all the bands have the same number of pixels, which is necessary as it requires fixed image sizes per channel. Depending on the specific application a higher or lower resolution may be used. Note this may slightly impact results when comparing merged to other compression methods. Except for this step, no further processing steps were applied to the bands.

3.1.2 Scenario

The scenario we used to demonstrate the feasibility of our approach is a constellation of three 6U CubeSats in the same orbit as Sentinel-2. Figure 3 gives a detailed overview of the parameters of the satellites and the constellation. We assume a Unibap iX10-100 device will be used for the training. Through profiling, we identified the time required to train on a batch on this device to be ten seconds. A detailed overview of the training algorithm is given in Section 2 and

[3] Sentinel-2 Products Specification Document. Available online at: "https://sentinel.esa.int/documents/247904/685211/sentinel-2-products-specification-document". Accessed: 2024-04-12

[4] https://github.com/ESA-PhiLab/PyRawS Accessed 2024-04-12

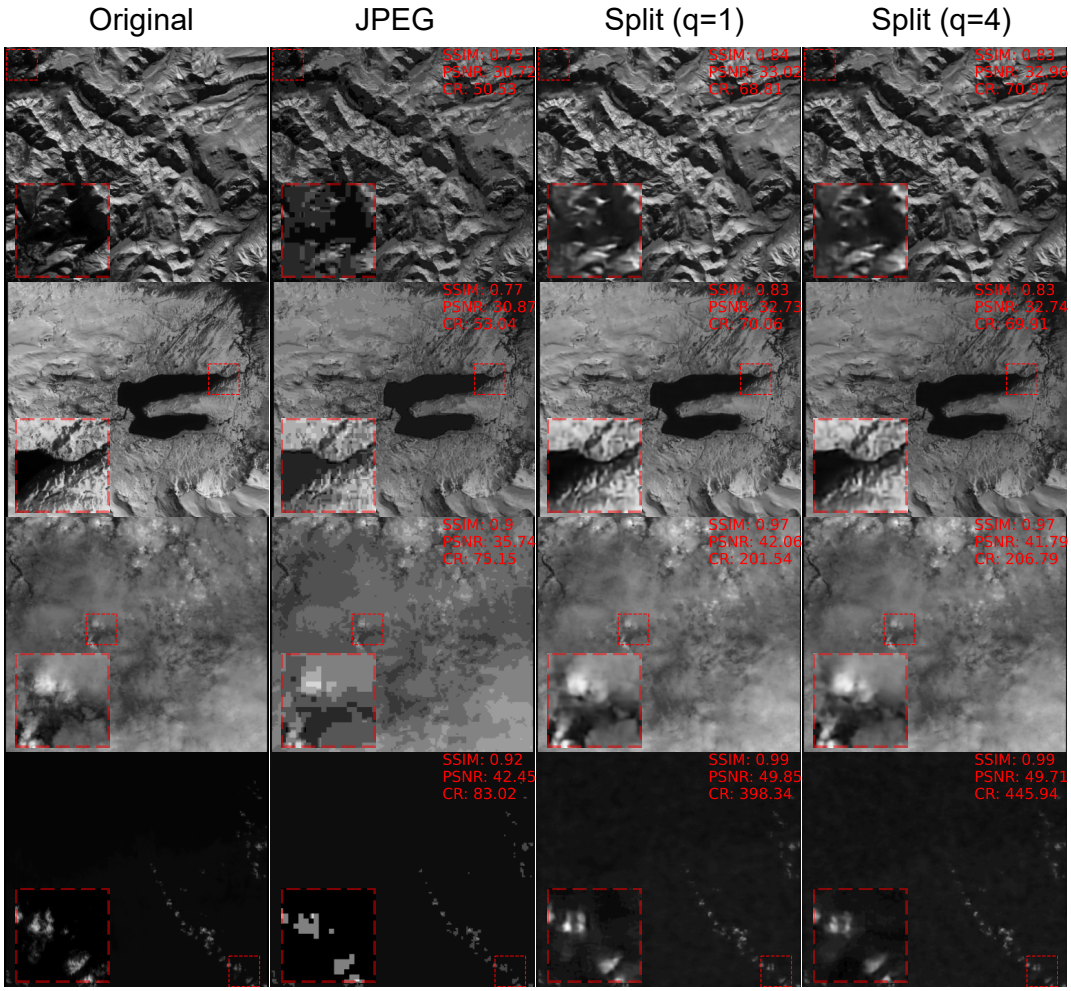

Figure 5. Exemplary compression results.

illustrated in Fig. 2. Most notably, the assumed satellites are smaller than the Sentinel-2 satellites to show that this approach can work on CubeSats.

3.2. Compression performance

Overall, we present results compared to the established JPEG standard regarding PSNR, SSIM, and BPP obtainable with LICOS. We also provide results in terms of compress ratio (CR), defined as the ratio between the compressed and the original file sizes.

To calculate JPEG results, we compressed the test set designed for the "split" model by using JPEG with quality factors [1,2,3,4,5,6,10]. The "merged" and "split" models were run on their respective test datasets, respectively with 13 and 1 input bands per image, with quality factors [1, 2, 4, 8]. Such quality factor values are supported by CompressAI for the autoencoder model used.

Fig. 4 we display a comparison with JPEG in terms of BPP/SSIM and BPP/PSNR values. Both the "merged" and "split" models achieve higher SSIM and PSNR values for lower BPP values than JPEG. Even with quality factor 1, JPEG cannot reach the same compression level as "merged" and "split" models, achieving a minimum BPP of 0.0985, with SSIM and PSNR values ranging from 0.585 to 0.892 and 29.5 to 37.3. Differently from JPEG, "merged" and "split" models achieve roughly constant SSIM, BPP, PSNR values for the different quality factors. Specifically, "split" models achieve average PSNR and SSIM values of 43.4 and 0.955 for an average BPP of 0.0383, while "merged" models feature average SNR and SSIM values of 37.7 and 0.950 for an average BPP of 0.00788. Notably, "merged" models achieve higher compression factors than split models at quality expenses, as demonstrated by the lower PSNR values.

Fig. 5 compares different Sentinel-2 raw bands after compression and decompression performed with different

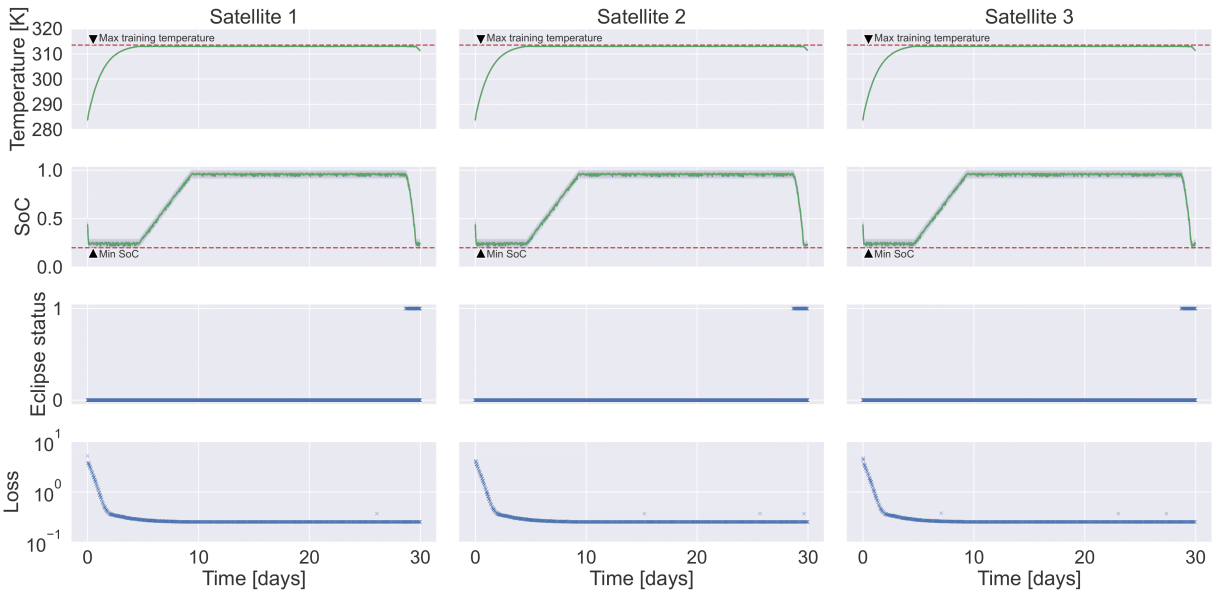

Figure 6. Temperature, state-of-charge, eclipse status, and Loss trends over time. Eclipse status is 1 and 0 respectively when the satellite is/is not in eclipse.

algorithms: JPEG (q=8), "split" model (q=1) and "split" model (q=4). Compression artefacts, PSNR, SSIM, and compression ratio values are shown for the different images and compression methods. Notably, the abundance of spatial features limits the compression factor that can be achieved. Indeed, "split" models offer slightly higher values of compression while ensuring higher quality as demonstrated by the higher values of PSNR, SSIM and lower number of artefacts. For the two images at the bottom, "split" models offer compression factors up to 5 times higher than JPEG while keeping higher quality after compression.

3.3. Operational constraints

Aside from the compression and image quality, we also investigate the operational aspects modeled with PASEOS. The internal temperature, state of charge (SoC), eclipse status and training loss trends are displayed in Fig. 6 for the different satellites in the constellation. At the beginning of our simulation, each satellite has a SoC of 50% and a temperature of 283.15 K.

The satellites are constantly exposed to solar radiative power, as they do not enter eclipse until day 28 of the orbit. This condition enhances the battery charging and increases the incoming heat flux, which adds to the heat dissipation from the internal sources during different activities. The current satellite design cannot dissipate the excess heat completely, resulting in a rise in the internal temperature that triggers interruptions in the training. For constant training and during eclipse, the power demand for the training activity also exceeds the power supply from the solar panels, leading to a decrease in the state of charge (SoC) of the battery.

Because of that, as shown in Fig. 6, at the beginning of the simulation, the SoC of the different satellites decreases from 0.5 to 0.2, forcing the satellite to switch to the "standby" activity until the battery has sufficient energy for the satellite to resume training or perform a model update. Specifically, during the initial five days, the SoC exhibits periodic fluctuations within the range of 0.2 and 0.29, whereas the satellites' temperatures rise steadily but remain below 313.5 K. (40° C). In this situation, the satellites are power-constrained since the low SoC triggers the "standby" activity and forces the training and model update to be stopped periodically.

Around Day 5, the satellite's internal temperature reaches 313.5 K, the maximum permissible temperature. Consequently, the "standby" activity is triggered to allow heat dissipation. Given the reduced power draw due to heat dissipation "standby", this results in the SoC increasing and stabilizing at approximately 0.95, with minor variations between 0.91 and 0.99. In this phase, satellites are constrained by their internal temperature.

Finally, between Days 28 and 29 satellites go into eclipse. This prevents the battery from being recharged while reducing the incoming heat flux. In this situation, both the SoC and the temperature decreases. A prolonged eclipse would make the satellite's batteries reach the limit of 0.2 SoC, making the satellite enter a power-constrained

state as at the beginning of the simulation.

Notably, since all the satellites exhibit similar SoC and temperature trends, all the training losses show a similar smooth decreasing trend. Therefore, differently from the behaviour showcased by Gome'z et al. [17], all the satellites provide a similar contribution to the model training in our scenario. Moreover, Figure 6 shows that the training loss for all the satellites dropped by more than 99% of the total delta already after Day 5. Therefore, we can reasonably infer that the training time can be substantially reduced with minimal performance loss, resulting in significant operational benefits. This possibility will be further explored in future works through a dedicated ablation study.

4. Discussion

4.1. Hardware limitations and satellite design

One limitation inherent to this approach is the need for suitable computing hardware. Especially the memory footprint of 5.472 GB (5218 MiB) is restrictive. Note that we only used the CPU of the Unibap iX10-100 device; on a GPU the training would be significantly faster. Thus, at this time, a state-of-the-art modern satellite processor is necessary. In the future, smaller, more efficient devices may also support this scenario. Running the training on a GPU would be most preferable and likely speed up training.

With the need for such a processor comes some design considerations that have to be made. The allocated power budget of 30W for the training is fairly high for a 6U CubeSat, but not unimaginable as there are ongoing studies, e.g., for a 100W 1U CubeSat deployable solar panel [27]. On the thermal management, one can already see from the PASEOS specifications in Fig. 3 that a fairly large emissive area to dissipate heat is assumed. This is likely necessary but not an insurmountable problem.

Similarly, the need for a large onboard mass memory is another possible limitation, especially if the selection of the previews of interest and the download of the corresponding useful data cannot be performed in a single pass with the ground station. This aspect is not investigated in this study. Indeed, during the training, only the model parameters are exchanged between the ground station and a satellite, and the operational life is not investigated after training.

4.2. Application to other mission scenarios

Although we explored a specific scenario in terms of constellation and data in this work, the presented approach should translate well to different types of data and applications. It can be employed for any imaging data, possibly also synthetic aperture radar, as it works on raw data. Further, it may be particularly beneficial in scenarios where it is not beforehand clear which images are of interest. This may, e.g., be the case in natural disaster detection. In these cases, one can send down the heavily compressed versions and even run potential detection algorithms to identify which images will be sent down with lossless compression [32].

4.3. Future work

This work presents the first results of the LICOS concept investigation. Results showcase that the proposed neural network model is a promising solution for transmitting heavily compressed previews, especially when the "split" is used. However, a proper thorough statistical evaluation including additional random seeds may help ensure the robustness of the approach. However, further investigations are needed to compare more thoroughly to other state-of-the-art compression methods for multispectral imagery, such as CCSDS 123-0-B-2 [22]. In that respect, we plan to consider additional metrics such as the mean absolute error and mean relative error that are relevant for remote sensing applications.

However, it is worth mentioning that we do not aim to replace existing standards or mission-specific solutions. Instead, we suggest complementing them with a low-overhead, lossy compression scheme that can allow transmission of heavily compressed versions of imagery that may otherwise be discarded or is not the primary objective. That way, data from areas underrepresented in datasets, such as oceans or other less often sampled areas, may also be analysed.

Furthermore, future investigations will focus on analysing the concept during the operational scenario after the initial model training. Indeed, as previously mentioned, current simulations focused only on the initial training phase, where the autoencoder model does not offer enough accuracy to be deployed. Because of that, aspects related to the satellite link budget, trade-offs in terms of onboard memory and number of ground stations available were not explored in the scope of this study.

5. Conclusion

This work presents the first results of the LICOS concept, which proposes the use of onboard-trained autoencoder models to deliver highly compressed previews to the ground mitigating the downlink bandwidth bottleneck. Two different configurations, "split" and "merged" models, were compared to JPEG standard on Sentinel-2 data demonstrating superior quality/compression trade-offs. Finally, an investigation of the main parameters affecting the training was investigated demonstrating the viability of training such models on properly designed 6U CubeSats. Results suggest the presented approach has the potential to mitigate the downlink bottleneck while providing the end users with the possibility to download original uncompressed images.

References

[1] Vinicius Alves de Oliveira, Marie Chabert, Thomas Oberlin, Charly Poulliat, Mickael Bruno, Christophe Latry, Mikael Carlavan, Simon Henrot, Frederic Falzon, and Roberto Camarero. Reduced-complexity end-to-end variational autoencoder for on board satellite image compression. *Remote Sensing*, 13(3):447, 2021. 2, 4

[2] Pascal Bacchus, Renaud Fraisse, Aline Roumy, and Christine Guillemot. Quasi lossless satellite image compression. In *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 1532–1535, 2022. 1, 2

[3] Pascal Bacchus, Renaud Fraisse, Aline Roumy, and Christine Guillemot. Quasi lossless satellite image compression. In *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 1532–1535. IEEE, 2022. 4

[4] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018. 1, 3, 4

[5] Jean Bégaint, Fabien Racapé, Simon Feltman, and Akshay Pushparaja. Compressai: a pytorch library and evaluation platform for end-to-end compression research. *arXiv preprint arXiv:2011.03029*, 2020. 3, 4

[6] Fredrik C Bruhn, Nandinbaatar Tsog, Fabian Kunkel, Oskar Flordal, and Ian Troxel. Enabling radiation tolerant heterogeneous gpu-based onboard data processing in space. *CEAS Space Journal*, 12(4):551–564, 2020. 2

[7] Aksel S Danielsen, Tor Arne Johansen, and Joseph L Garrett. Self-organizing maps for clustering hyperspectral images onboard a cubesat. *Remote Sensing*, 13(20):4174, 2021. 4

[8] Vinicius Alves de Oliveira, Marie Chabert, Thomas Oberlin, Charly Poulliat, Mickael Bruno, Christophe Latry, Mikael Carlavan, Simon Henrot, Frederic Falzon, and Roberto Camarero. Satellite image compression and denoising with neural networks. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. 4

[9] Roberto Del Prete, Nicolas Longepe, Federico Serva, Olivier Colin, and Alix De Beusscher. Thraws [data set], 2023. Available online at: https://doi.org/10.5281/zenodo.7908728. 5

[10] Roberto Del Prete, Gabriele Meoni, Nicolas Longépé, Maria Daniela Graziano, and Alfredo Renga. First results of vessel detection with onboard processing of sentinel-2 raw data by deep learning. In *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*, pages 6262–6265. IEEE, 2023. 1, 5

[11] Maria Pia Del Rosso, Alessandro Sebastianelli, Dario Spiller, Pierre Philippe Mathieu, and Silvia Liberata Ullo. On-board volcanic eruption detection through cnns and satellite multispectral imagery. *Remote Sensing*, 13(17):3479, 2021. 1

[12] Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. Sentinel-2: Esa's optical high-resolution mission for gmes operational services. *Remote sensing of Environment*, 120:25–36, 2012. 2, 4

[13] Yann Dubois, Benjamin Bloem-Reddy, Karen Ullrich, and Chris J Maddison. Lossy compression for lossless prediction. *Advances in Neural Information Processing Systems*, 34:14014–14028, 2021. 3

[14] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 3

[15] Gianluca Furano, Gabriele Meoni, Aubrey Dunne, David Moloney, Veronique Ferlet-Cavrois, Antonis Tavoularis, Jonathan Byrne, Léonie Buckley, Mihalis Psarakis, Kay-Obbe Voss, and Luca Fanucci. Towards the use of artificial intelligence on the edge in space systems: Challenges and opportunities. *IEEE Aerospace and Electronic Systems Magazine*, 35(12):44–56, 2020. 1, 2

[16] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batič, Léonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefele, Nathan Vercruyssen, 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*, 60:1–14, 2021. 1, 2, 4

[17] Pablo Gómez, Johan Östman, Vinutha Magal Shreenath, and Gabriele Meoni. Paseos simulates the environment for operating multiple spacecraft. *arXiv preprint arXiv:2302.02659*, 2023. 2, 3, 8

[18] Giorgia Guerrisi, Fabio Del Frate, and Giovanni Schiavon. Artificial intelligence based on-board image compression for the ϕ-sat-2 mission. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2023. 1, 2

[19] Dailan He, Ziming Yang, Weikun Peng, Rui Ma, Hongwei Qin, and Yan Wang. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5718–5727, 2022. 3

[20] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In *IGARSS 2018-2018 IEEE international geoscience and remote sensing symposium*, pages 204–207. IEEE, 2018. 2

[21] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. 2

[22] Miguel Hernández-Cabronero, Aaron B Kiely, Matthew Klimesh, Ian Blanes, Jonathan Ligo, Enrico Magli, and Joan Serra-Sagrista. The ccsds 123.0-b-2 "low-complexity lossless and near-lossless multispectral and hyperspectral image compression" standard: A comprehensive review. *IEEE Geoscience and Remote Sensing Magazine*, 9(4):102–119, 2021. 1, 8

[23] Dario Izzo, Gabriele Meoni, Pablo Gómez, Domink Dold, and Alexander Zoechbauer. Selected trends in artifi-

cial intelligence for space applications. *arXiv preprint arXiv:2212.06662*, 2022. 4

[24] Gonéri Le Cozannet, M Kervyn, S Russo, Chinwe Ifejika Speranza, P Ferrier, Michael Foumelis, Teodolina Lopez, and H Modaressi. Space-based earth observations for disaster risk management. *Surveys in geophysics*, 41:1209–1235, 2020. 1

[25] Bho Matthiesen, Nasrin Razmi, Israel Leyva-Mayorga, Armin Dekorsy, and Petar Popovski. Federated learning in satellite constellations. *IEEE Network*, pages 1–16, 2023. 2

[26] Johan Östman, Pablo Gomez, Vinutha Magal Shreenath, and Gabriele Meoni. Decentralised semi-supervised onboard learning for scene classification in low-earth orbit. *Proceedings of the 14th IAA Symposium on Small Satellites for Earth Observation, SSEO*, 2023. 3, 4

[27] Antonio Pedivellano, Thomas Sinn, Ambre Raharijaona, Michael Kringer, Joram Gruber, Joachim Schmidt, Thomas Lund, Alexander Titz, Diego Garcia, Daria Stepanova, et al. Prototyping and engineering model test campaign of the 100w 1u powercube deployable solar array. In *AIAA SCITECH 2023 Forum*, page 2696, 2023. 8

[28] Mihir Prakash, Steven Ramage, Argyro Kavvada, and Seth Goodman. Open earth observations for sustainable urban development. *Remote sensing*, 12(10):1646, 2020. 1

[29] Nasrin Razmi, Bho Matthiesen, Armin Dekorsy, and Petar Popovski. Ground-assisted federated learning in leo satellite constellations. *IEEE Wireless Communications Letters*, 11 (4):717–721, 2022. 1

[30] Nasrin Razmi, Bho Matthiesen, Armin Dekorsy, and Petar Popovski. On-board federated learning for satellite clusters with inter-satellite links. *IEEE Transactions on Communications*, 2024. 2

[31] John A Richards, John A Richards, et al. *Remote sensing digital image analysis*. Springer, 2022. 5

[32] Vít Růžička, Anna Vaughan, Daniele De Martini, James Fulton, Valentina Salvatelli, Chris Bridges, Gonzalo Mateo-Garcia, and Valentina Zantedeschi. RaVÆn: unsupervised change detection of extreme events using ML on-board satellites. *Scientific Reports*, 12(1):16939, 2022. Number: 1 Publisher: Nature Publishing Group. 1, 8

[33] Susan L Ustin and Elizabeth M Middleton. Current and near-term advances in earth observation for ecological applications. *Ecological Processes*, 10:1–57, 2021. 1