SSC23-XIII-08

Adaptive On-board Signal Compression for SAR Using Machine Learning Methods

Craig Hay, Lucy Donnell, Charlotte Crawshaw, Murray Ireland
Craft Prospect Ltd
Suite 12 Fairfield, 1048 Govan Road, Glasgow, United Kingdom, G51 4XS; +44 (0) 7421 994 712
craig.hay@craftprospect.com

Mehrdad Yaghoobi
School of Engineering, University of Edinburgh
204 AGB Building, EH9 3FG, United Kingdom; +44 131 650 7185
m.yaghoobi-vaighan@ed.ac.uk

ABSTRACT

Satellites with synthetic aperture radar (SAR) payloads are growing in popularity, with a number of new institutional missions and commercial constellations launched or in planning. As an active instrument operating in the microwave region of the electromagnetic spectrum, SAR provides a number of unique advantages over passive optical instruments, in that it can image in all weather conditions and at night. This allows dense time-series to be built up over areas of interest, that are useful in a variety of Earth observation applications. The polarisation and phase information that can be captured also allows for unique applications not possible in optical frequencies.

The data volume of SAR captures is growing due to developments in modern high-resolution multi-modal SAR. Instruments with higher spatial resolution, wider swaths, multiple beams, multiple frequencies and more polarization channels are being launched. Miniaturization and the deployment of SAR constellations is bringing improved revisit times. All of these developments drive an increase in the operational cost due to the increase in data downlink required. These factors will make on-board data compression more crucial to overall system performance, especially in large scale constellations.

The current deployed state-of-the-art of on-board compression in SAR space-borne payloads is Block Adaptive Quantization (BAQ) and variations such as Flexible BAQ, Entropy Constrained BAQ and Flexible Dynamic BAQ. Craft Prospect is working on an evolution of these techniques where machine learning will be used to identify signals based on dynamics and features of the received signal, with this edge processing allowing the tagging of raw data. These tags can then be used to better adjust the compression parameters to fit the local optimum in the acquired data.

We present the results of a survey of available raw SAR data which was used to inform a selection of applications and frequencies for further study. Following this, we present a comparison of a number of SAR compression algorithms downselected using trade-off metrics such as the bands/applications they can be applied to and various complexity measures. We then show an assessment of AI/ML feasibility and capabilities, with the improvements assessed on mission examples characterised by the SAR modes and architecture for specific SAR applications. Finally, future hardware feasibility and capability is assessed, targeting a Smallsat SAR mission, with a high level roadmap developed to progress the concept toward this goal.

INTRODUCTION

Synthetic aperture radar (SAR) is a sensing modality that is of growing interest in the Smallsat domain. From the first civil mission, Seasat[1] in 1979, until recently, SAR has been the preserve of large institutional or government backed missions. In the recent class of ESA Copernicus scientific missions, the pair of Sentinel-1 satellites has made C-band SAR with consistent global coverage freely available for all, spurring innovation in downstream applications of SAR data.

On the platform side, new developments in semiconductor materials (e.g. Gallium Nitride) have made compact high-power RF amplifiers possible and allowed the miniaturization of instrument technology. This has enabled a multitude of commercial companies (e.g. SSTL, ICEYE, Synspective, Capella, Umbra) to enter the SAR observation market, offering more frequent coverage, tasking capabilities and high resolution spotlight imaging of regions-of-interest. Not only are these Smallsat operators augmenting existing services, but also providing new capabilities. For example, ICEYE "Dwell Mode" long staring spotlight[2] mode, or a consortium led by Umbra looking at bistatic applications under the DARPA funded Distributed Radar Image Formation Technology (DRIFT) program.[3]

The attraction of SAR in comparison to electro-optical imaging technologies is the ability to image in all-weather and all-lighting conditions which allows customers to build dense time series over areas of interest. The deployment of constellations of spaceborne SAR platforms enhances this capability. SAR also offers some unique applications based on the physics of the received signal. Due to the ability to measure both magnitude and phase of signals at microwave frequencies of the EM spectrum, processing techniques can be applied to make measurements that are not possible with the radiometric measurements of electro-optical devices. Space-based SAR interferometry is an application that has seen widespread use in fields like seismology, disaster management and topographic measurement. The active nature of the SAR instrument also allows control of the polarization of emitted signals, allowing the recording of the polarization behaviour of scatterers on the ground. From these measurements we can extract information such as land-use classification, agricultural monitoring and forest biomass estimation.

In summary, spaceborne SAR system architectures are undergoing a step change in that it is becoming possible to deploy large constellations of satellites as opposed to a single instrument or pair of satellites. On larger SAR missions, the trend is toward electronically steerable active phased array antenna and multi-beam/multi-frequency instruments. The sheer amount of data generated by these instruments presents a systems-level challenge for space-ground downlink that requires novel approaches to overcome. It likely will not be possible to simply increase downlink capacity due to satellite power and bandwidth constraints and the limited spectrum available due to radio license constraints. It is from these system-level motivations that the desire for more effective on-board compression stem.

SYSTEM ARCHITECTURES

This work on more intelligent on-board data compression is planned to be one of the first steps in a program of work to enable more efficient use of SAR space assets.

Changing SAR satellite system architectures are moving from single and/or pairs of satellite systems to the deployment of large constellations. This brings improvements in constellation robustness, reliability and resilience as well as performance improvements such as greater tasking and a reduction of latency from initial tasking by customer to delivery of data product.

For example, an architecture using six polar orbital planes, each containing six satellites equal phases apart, provides something approaching global coverage, means that there is a pass approximately every fifteen minutes over an area. With timely analysis of the data collected by a single satellite observation, this provides time for subsequent satellites to image the same area with varying parameters. The simplest example of constellation tasking in this manner would be to command the subsequent satellite to acquire another image of the same region.

SAR platforms are typically heavily duty cycled, this is especially apparent on Smallsat platforms, with technical details provided by Synspective at past Smallsat conferences[4] quoting typical figures of 5% or 10% per orbit, or 25% in their case. Adding more satellites to a constellation can mitigate this issue, though only up to a point because as well as providing new capabilities, these multi-satellite systems also impose greater costs on supporting infrastructure such as ground station facilities which have finite capacity.

The hard duty cycling imposed on each individual platform is due to the active nature of the radar imaging sensor. To get the required range return from the side-looking instrument, 1-2kW power is transmitted. With power amplifiers of around 50% efficiency generating these high-power RF signals, this means 3-4kW of heat being generated on-board. On these small, simple platforms with no active cooling mechanisms and small power budgets (on the order of 500W solar cell capacity) this is both a challenge in platform power generation and a thermodynamic problem in heat dissipation. Thermal deformations, especially in antenna, can prevent accurate images from being acquired, so it is a sensitive thermal environment.

With instruments acquiring data at rates of Gbps when observing, dealing with this amount of data and downlinking it through the space-ground link is also a challenge. Any effort to reduce this on-board will provide increases in system capacity. The motivations for better on-board compression are therefore:

- The low duty cycle of SAR instruments
- Lowering on-board storage requirements or increasing amount that can be stored on-board
- Lowering the amount of data that needs to be downlinked

Further higher-level operations that on-board processing can enable are:

- Prioritising data for downlink
- Tasking satellites with extracted information
- Change detection based on compression performance

MACHINE LEARNING FOR COMPRESSION

Artificial intelligence (AI) and machine learning (ML) solutions have been successfully harnessed to enable on-board processing of optical satellite data in applications such as on-board cloud detection.[5] These solutions reduce data bottlenecks that result from the rapidly growing amount of data collected on-board satellites and the limited downlink bandwidths available.[6] Where high value and time critical data can be prioritised by real-time intelligent decision making, the bottleneck problem is overcome.

The operational concept of AI enabled onboard data processing can be harnessed for various Earth observation applications including event detection, prediction and monitoring; persistent monitoring; damage assessment; global and localised mapping; object detection and tracking, and scientific measurement.[7] Whilst in these applications trained models are deployed and used for *information extraction* there is much literature to suggest that trained models are also effective *lossy compressors*,[8] which is the task required in on-board processing of raw SAR data.

A traditional compression algorithm will be designed to remove duplicate or redundant information before encoding in the fewest possible bits, with the algorithm designer exploiting prior knowledge about the source to be encoded during the design process. In machine learning, as a data driven discipline, the model is shown input and output data and through a process of iterative optimisation will converge on a model that minimizes some designed loss function used in training. For this process to be successful, sufficient amounts of data are required to train.

AVAILABLE DATA

To study the possible improvements in machine learning driven compression, datasets are required for training and testing the machine learning models and for comparisons with traditional algorithms. A dataset survey was conducted in the initial stages of the work.

When surveying datasets it is important to make the distinction between Data, Dataset and Machine Learning Ready Dataset. "Data" requires no definition. A "Dataset" is a bounded, finite collection of data that has some rationale behind it for being grouped together, a curated collection of data. A "Machine Learning Ready Dataset" is a Dataset that contains everything required to begin training a machine learning algorithm in some chosen task. This requires a Dataset to have additional properties such as being formatted for easy ingestion into machine learning training pipelines, possibly with labelled or derived features for supervised learning techniques to be applied. There should also be enough data to allow the partitioning of the dataset into training/test/validation sets, and the data should have enough diversity that it adequately covers the problem space.

The availability of such ML-ready datasets vastly reduces the effort required to get up and running with machine learning experimentation and research, and likely spurs innovation in the application area covered by the dataset (see ImageNet[9] for an example of this). Specifically constructed benchmark datasets can serve a purpose in allowing fair comparisons in compression performance between different works using the same dataset, see for example the CCSDS Reference Image Set.[10]

The conclusions of the survey found that although there is some raw SAR data available, there are no publicly available raw SAR datasets, and no ML-ready datasets. This situation with the lack of raw SAR datasets can be contrasted to the availability of SAR Single Look Complex (SLC) and Ground Range Detected (GRD) products, where datasets are available, and even in a ML-ready dataset format for certain applications. These existing datasets are mainly targeted at downstream applications such as:

- Object detection such as ship detection, classification and identification
- Fusion of SAR and optical imagery for cloud in-painting

- Speckle filtering using deep learning

To do any higher-level processing on-board, such as object detection, significant processing is required initially to get from raw SAR data to a Level-1 GRD product, which has been detected, multi-looked and projected to ground range using an Earth ellipsoid model. This appears to be a challenge for the current state-of-the-art in avionics and an area of current research.[11]

After conducting a survey on available sources of raw SAR data, Sentinel-1 was chosen as the source. Raw Sentinel-1 data is freely available on Copernicus Open Access Hub.[12] It should be noted that, as a C-band satellite, this has been chosen for data availability reasons. X-band is also of interest as it is common in Smallsat applications, but the raw data availability is not there. Selected raw SAR data in C-band is also freely available from ESA for older SAR missions like ERS-1, ERS-2 and ENVISAT.

Sentinel-1 raw data has already been preprocessed with a lossy compression applied on-board, the FD-BAQ algorithm. The ideal for compression experiments is "BAQ bypassed" data, that is typically only downlinked for commissioning and calibration reasons as it has no compression applied so is downlinked at the full instrument ADC sample depth (e.g. 8-bits per sample). For ML-driven compression work, a dataset of {*image content label, raw data, focused image*} triples would be ideal for exploring opportunities in ML-assisted SAR data compression.

METHODOLOGY

A methodology has been adapted to assess the benefits of further on-board data compression. This is shown in Figure 1.

From the decision to use Sentinel-1 data exclusively, a pipeline then has been developed to ingest raw data from Copernicus Open Access Hub. The pipeline first processes the raw data to a machine-learning ready format, trains machine learning algorithms, compresses/decompresses using these algorithms or traditional compression algorithms, then quantitatively compares results using a series of data metrics in the raw data domain, and image metrics after the results have been focused into SAR images. The focusing is based on the SENTINEL1DECODER[13] library.

Areas acquired in Stripmap Mode (SM) will be preferred, as the focusing and image interpretation is a simpler task in this mode when working with the raw data file format. Sentinel-1 usually only acquires Stripmap over small island targets, where the wider swath of Interferometric Wave (IW) mode is not required. There are some other scenes where Stripmap Mode is regularly available, which are likely test sites.

The SENTINEL1DECODER library examples, and previous work on raw SAR data compression,[14] use Sao Paolo, Brazil as the Stripmap Mode target which looks to provide a good mixture of busy shipping port, coastal area, urban areas and forest canopy. This early work also used only 5 Sentinel-1 scenes for training of variational autoencoders for compression. If there is a need identified that can only be satisfied by scenes acquired in IW mode, the swath number metadata will be used to select and focus the subswaths separately, which will then be chipped into small blocks for training. This will avoid having to merge subswaths, which would add to the complexity. Scenes are very large, containing many cells (e.g. 29934 x 19950) that are often chipped into smaller images to create very long sequences of smaller images for training. Raw data will be downloaded and decoded offline with SENTINEL1DECODER, with the IQ data decoded and stored in .npy files for easier manipulation into machine learning ready formats, and the metadata stored in .csv or equivalent.

The statistical analyses and image domain metrics will be calculated offline for the scenes, and the results stored. This work presents computational challenges due to the large amounts of raw data (gigabytes per scene), high computational effort for focusing, and the data processing chain needed to turn data into machine-learning ready dataset. Working on a very small subset of Sentinel-1 scenes (5 to 10) will ease this.

This software pipeline can be adapted to X-band data in future by modifying the "Read uncompressed raw data" and "SAR focusing" blocks, to better target data sources and frequency that is heavily used in Smallsat commercial SAR.

Standard statistical measures will be calculated in the data domain.[15] These are simple to calculate and provide some measure of how the data has changed before and after compression. Although the compression is lossy, there should be no significant change in the overall statistics of the data after the compression algorithm has been applied. There are well-known statistical models of SAR radar echoes in the data domain.[16] The model consists of:

- Complex data, sampled at 8-bits/sample
- Zero-mean, circular Gaussian distribution of complex samples

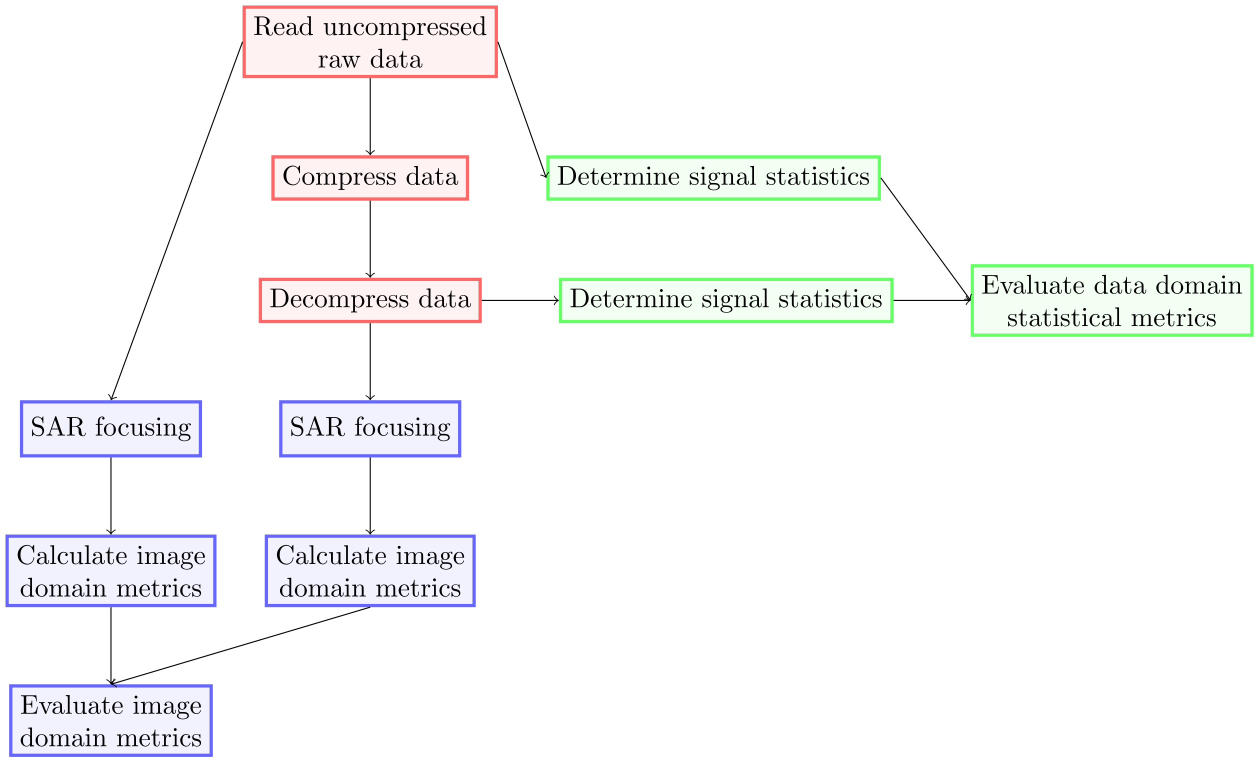

Figure 1: Algorithm assessment methodology (image domain in blue, data domain in green)

- Small amount of saturation in analog-to-digital converter (ADC)
- Low correlation between I and Q channels
- Low intersample correlation in range and azimuth
- Slowly changing variance in slant range and azimuth

The I and Q signals in rectangular form follow Gaussian distributions with mean zero and variance dependent on the data. The I and Q signals are uncorrelated. This holds particularly well for homogeneous scattering regions. IQ signals are complex valued measurements so some measures need to be assessed for I and Q separately. The rectangular form can be converted to polar form to achieve magnitude/phase representation of the signal. This changes the statistics from Gaussian to a Rayleigh distribution for the magnitude and a uniform distribution for the phase.

The metrics used to monitor compression performance in the data domain will be:

- Dynamic range
- Mean (magnitude and phase)
- Standard deviation (magnitude and phase)
- Skewness (magnitude and phase)
- Kurtosis (magnitude and phase)
- Entropy (magnitude and phase)

Further details on metrics can be found in this survey.[15]

These measures in the data domain only check whether the compression algorithm has drastically changed the data or not, which is not sufficient for actual performance analysis. The perils of relying on only summary statistics are shown in Anscombe's Quartet,[17] which has examples of four constructed datasets with clearly very different underlying trends when visualized, but with the same summary statistics. Visualization of data in the image domain is important. Raw SAR data is especially difficult to visualize due to the data properties quoted in the above model. Therefore, part of the assessment must include focusing the raw data into SAR images.

As well as conducting visual inspections of focused images, a set of image metrics have been employed to assess compression performance. To capture the ability to preserve point targets in the image, Peak to Sidelobe Ratio (PSLR) measurements have been employed as a worst-case measure of SAR ability for identifying a weak target from a nearby strong target. To conduct this analysis, the ideal for this is corner reflectors or transponders in the images, but as this study is using real data we have only real targets to choose from. To capture overall error in the amplitude image, mean squared error (MSE) will be used. The dynamic range of the image will also be assessed in the image domain.

The preservation of phase will be particularly important as accurate phase preservation is a requirement for interferometric SAR. In future works, the data produced by any developed compression algorithm will require testing in these derived applications of SAR data such as interferometry, where phase preservation is important and phase error leads to distance measurement error).

STANDARD ALGORITHMS

The standard methods of on-board SAR data compression rely upon the Gaussian nature of the signal. These assumptions mean a simple algorithm can be developed for compression, Block Adaptive Quantization (BAQ). The BAQ algorithm was first presented for implementation on NASA's Magellan probe[18] which began design in the late 1970s and was finally launched on its mission to radar image Venus in 1989. The simplicity of the algorithm for hardware implementation was paramount, though as it has proven to be quite effective at compressing this type of data it has been extended and improved upon in the years since. The BAQ algorithm is detailed in Figure 2.

BAQ aims to reduce the data volume while preserving important information within acceptable tolerances. The BAQ algorithm makes several assumptions during the compression process:

Statistical Independence BAQ assumes that the real (I) and imaginary (Q) components of the complex SAR data are statistically independent. This assumption allows for separate quantization and compression of the I and Q components.

Dynamic Range BAQ assumes that the dynamic range of the SAR data is approximately uniform across a block. It assumes that the backscattering characteristics and clutter statistics remain relatively consistent. This assumption allows for a uniform allocation of quantization bits across the entire data, ensuring that no specific regions are favored or compromised during compression.

Limited Visual Impact BAQ aims to reduce the data size while maintaining a visually acceptable image quality. The algorithm assumes that certain compression artifacts and information loss can be tolerated without significantly degrading the interpretability of the SAR image. The specific tolerances for compression-induced artifacts may vary depending on the application and end-user requirements.

BAQ operates by calculating the variance of a block and using this summary statistic to place the thresholds of the quantizer optimally (4 thresholds for 2-bit BAQ, 8 for 3-bit and so on). The optimal placement for different numbers of threshold in uniform quantizers given the signal variance were proven and calculated in Max's seminal paper[19] in 1960. Using these values in both encoder and decoder, and transmitting block variance as side information, the decoding process is simply the reverse of the encoding. Even with this specification there are a number of trade-offs required in the design of a BAQ implementation. The sizing of blocks is important, too large and the assumption of stationary variance across the block breaks. Too small and the amount of side information (block standard deviations, typically encoded in 8-bits) becomes too much of an overhead.

As the BAQ method uses a fixed codeword length, for instruments with multiple modes a flexible variant named Flexible BAQ[16] (FBAQ) was developed, which allowed ground selection of the bitrate (e.g. 8 to 4, 8 to 3 or 8 to 2-bit). This was flown on ENVISAT. This added flexibility was important for a (relatively) modern SAR sensor such as ASAR on ENVISAT, as the different operating modes of the sensor and applications of the SAR images have different image quality and data volume requirements.

It is known from image compression work that applying transform techniques before coding can result in improvements in compression performance. In images this is because the spatially correlated information in the image domain is transformed to the frequency domain with Fast Fourier Transform (FFT) or Discrete Cosine Transform (DCT), or wavelet domain with different wavelet transforms. These encode the data as a series of coefficients in the new basis, these coefficients can be quantized separately or zeroed entirely in the coding step (e.g. removing very high frequency information that is an acceptable loss when transformed back into the image domain). This technique has been applied to BAQ with the FFT-BAQ algorithm which performs the transform to the frequency domain before encoding using standard BAQ.

Entropy coding is another technique from data compression that has been applied to BAQ, with Entropy Constrained BAQ (EC-BAQ). Entropy coding is a technique where the source probabilities of the

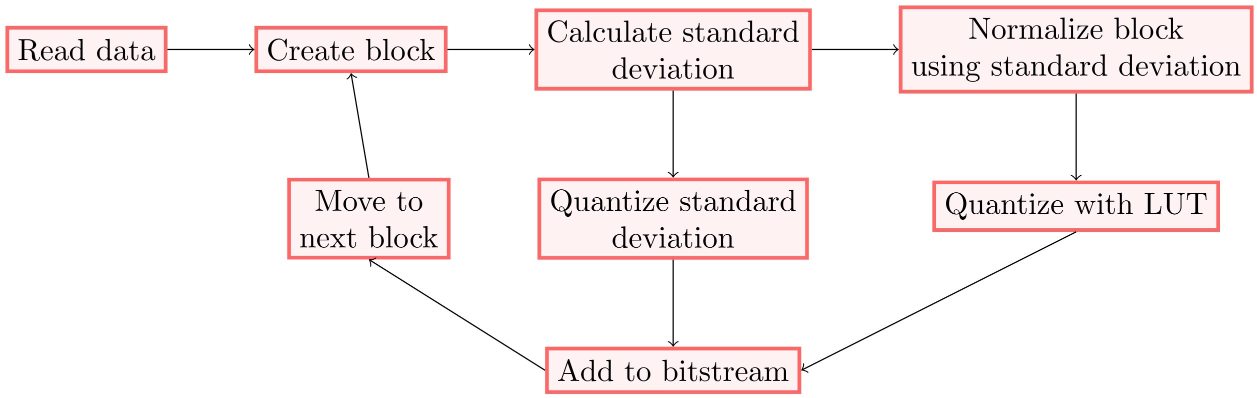

Figure 2: BAQ algorithm

alphabet of data to be encoded is considered in the design of the code. The more probable symbols in the alphabet are allocated shorter codewords and the less probable symbols longer codewords. There are two possible approaches to variable length coding of quantizer outputs. Either keep the design of the quantizer the same and simply entropy code the output of the quantizer, or move quantizer decision boundaries (introducing non-uniform decision boundaries) by taking into account how the selection of these will affect rate. In EC-BAQ the probabilities of the data in a block are used to design the steps of the quantizer that are used per block. The sizes of the decision boundaries are adjusted so that they are non-uniform. This allows the encoding of non-integer numbers of bits per block, as opposed to standard BAQ where each block will be a multiple of 2, 3 or 4 bits in size depending on the chosen codeword length.

A current state-of-the-art algorithm is Flexible Dynamic BAQ (FDBAQ), as flown on Sentinel-1. In a purely homogeneous scene FDBAQ performs similarly to fixed-rate BAQ.[20] The quantizer thresholds are set by taking into consideration the thermal noise measurements made by the system in flight. This measurement is used to choose a quantizer out of a set of quantizers based on the local signal-to-thermal-noise ratio, which is estimated blockwise on the data.

MACHINE LEARNING FEASIBILITY

To explore the feasibility of machine learning approaches, two paradigms were explored. One pure learning approach, and another enhancing aspects of existing algorithms by building on their design with a machine learning element.

Autoencoders are deep artificial neural networks that aim to reproduce their input at their output layer. When trained successfully, they are more effective than traditional compression techniques such as Principal Component Analysis for dimensionality reduction.[21]

Variational Autoencoders show great promise in image data compression in unsupervised learning regimes. Variational inference is an essential statistical learning technique where a distribution (such as Gaussian) is selected and its parameters adjusted to best match a target distribution, even when the target distribution is not exactly known. Autoencoders consist of an encoder and decoder as artificial neural networks, which are trained to learn the best encoding and decoding scheme by performing iterative optimisation. The output of the autoencoder network is compared with the input data to measure the error, which is then back propagated through the model to update the network weights in a way that best reduces the error. Variational autoencoders are those which have an encoding distribution that is regularised during training. Regularisation is achieved by a modification of the encoding and decoding process where the input is not encoded as a single point of data, but rather as a distribution over the latent (target) space.

An approach explored for the raw SAR data compression application is using Vector Quantized Variational Autoencoders (VQ-VAE), where discrete (rather than continuous) latent variables are used and distributions are categorical.[22] It is important to note that the raw SAR data collected onboard is not a suitable end product for users on the ground. Typically, analysis of SAR products downstream follows necessary radiometric calibration and focusing of the data. Successful on-board compression of SAR data means enabling efficient downlink of complex data with phase and amplitude preserved. A high compression ratio of the raw data must be achieved, while still attaining a high-quality end product in the image domain.[14]

In our first approach we explore various VQ-VAE based model architectures, parameters and performance metrics to reach these goals. We benchmark the performance against the state-of-the-art standard algorithms and evaluate the feasibility of AI as a solution to on-board SAR data compression. Within the model learning and evaluation process, the input is raw SAR data and we aim for the highest quality approximation, with maximum compression ratio, as output. Outside of this process, the model output is further evaluated and verified against Single Look Complex (SLC) data to assess the quality of the end product. As part of the feasibility evaluation, we also consider the assurance of the autonomous on-board compression component to promote trust. This is done through applying the Assurance of Machine Learning in Autonomous Systems[23] (AMLAS) approach throughout development of the VQ-VAE solutions.

As the second approach, we revisit the standard BAQ approach and develop an ML model adaptor to improve the overall performance of the method using the underlying structure of training data corpus. This approach is important due to providing statistical adaptivity to the standard methods, using more complex data structure/statistics. The ML based approach can challenge the current performance due to exploiting the dependency of data samples, which is not currently exploited when an i.i.d. distribution is assumed to derive the optimal setting for the BAQ method. Such a data inter-dependency has enabled high imagery data compression using deep neural networks,[24] i.e. a special and hugely successful form of ML models, see this[25] for a relevant approach and this survey[8] on the state of the art methods.

HARDWARE CONSIDERATIONS

Finally, future hardware feasibility and capability is assessed, targeting a Smallsat SAR mission, with a high level roadmap developed to progress the concept toward this goal.

The hardware considerations must be examined from a number of angles. Space environmental factors play a large part in device selection, as do device size, weight and power (SWaP) requirements and how they fit at the system level, and also the ability of the device to execute the data processing at the required throughput. Space environment considerations can dictate the choice of radiation-hardened by design (RHBD), radiation tolerant (RT) or commercial off-the-shelf (COTS) devices for integration in space platform architecture.[26] This choice will also be influenced by data processing throughput requirements, platform power budget and total mission cost and attitude to risk. RT or COTS devices typically offer superior performance per Watt when compared against their RHBD counterparts. There is a typical radiation hardness vs capability trade-off due to the lessened physical effects of Single Event Upset (SEU) from radiation on devices manufactured with larger process nodes (e.g. 250nm vs 8nm). Newer COTS devices on smaller process nodes typically come with physical size, silicon architectural improvements, and most importantly performance per Watt increase.

The latest COTS technologies adapted for the space environment provide the ability to hardware accelerate machine learning workloads on edge devices. Devices like the Xilinx Ultrascale+ MPSoC and Intel Myriad X provide these capabilities in packages of low SWaP suitable for Smallsat power budgets. Just as importantly, device specific tools are available for on-boarding trained models such as Intel OpenVINO and Xilinx Vitis AI. The tools for developing on COTS are typically more mature and have seen greater investment due to their large user base. A consideration of the tools and their capabilities will feed into algorithm development, as this work is being done whilst looking ahead to future hardware deployment of the developed algorithm.

The selection of specific devices for a future demonstrator will be impacted by the processing system architecture when configured from the reference data processing architecture. Where a selected device has native fault-tolerance (i.e. is RHBD and has integrated error-checking), less system-level fault tolerance is required. Where a less mature device, such as the Myriad X VPU, is selected, increased system-level fault tolerance and mitigation is required, such as more capable fault detection, isolation and recovery (FDIR) and possibly component duplication or triplication.

In general, for missions where ML-based onboard processing is used for information extraction it can be seen as an enhancement to functionality rather than critical to the system. When used for data compression, the functionality is likely to be a critical system requirement. Computing devices used in institutional missions have tight radiation requirements, usually fulfilled by RT/RHBD devices and are space qualified through a high level of product assurance, such as the QML Standard. In "newspace" missions, radiation requirements are looser due to a higher risk acceptance for the mission and usually shorter mission lifetime. This means COTS devices are often used to exploit their advancements in compute power, ease of use and lower cost that they offer over RT/RHBD devices. By defining mission and subsystem attitude to risk, we can then differentiate among devices and make device choices based on the criticality of the functions. Reliability, robustness and resilience can then be achieved in the most effective way either through RHBD devices, or device-level redundancy, or satellite (system) level redundancy in the case of a many satellite constellation, as part of a mission at scale.

References

[1] Rolando L. Jordan. The Seasat-A synthetic aperture radar system. *IEEE Journal of Oceanic Engineering*, 5(2):154–164, 1980.

[2] Vladimir Ignatenko, Matthew Nottingham, Andrea Radius, Leszek Lamentowski, and Darren Muff. ICEYE Microsatellite SAR Constellation Status Update: Long Dwell Spotlight and Wide Swath Imaging Modes. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pages 1493–1496, 2021.

[3] Umbra. DARPA Selects Umbra for Their DRIFT Program, May 2023. Accessed 06/06/2023. Available at https://umbra.space/blog/darpa-selects-umbra-for-their-drift-program.

[4] Hirobumi Saito, Kosei Ishimura, Jiro Hirokawa, Takashi Tomura, Koichi Ijichi, and Koji Tanaka. Technical Challenges for Small SAR Satellites with High Performance. In *Proceedings of the AIAA/USU Conference on Small Satellites*, 2021.

[5] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batic, Leonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefele, Nathan Vercruyssen, Gianluca Furano, Massimiliano Pastena, and Josef Aschbacher. The -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, 2022.

[6] 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.

[7] Murray Ireland, Craig Hay, Doug McNeil, Derek O'Callaghan, Sheila McBreen, and Roberto Camarero. Applications and Enabling Technologies for On-Board Processing and Information Extraction: Trends and Needs. In *European Workshop on On-Board Data Processing (OBDP2021)*, June 2021.

[8] Siwei Ma, Xinfeng Zhang, Chuanmin Jia, Zhenghui Zhao, Shiqi Wang, and Shanshe Wang. Image and Video Compression With Neural Networks: A Review. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(6):1683–1698, June 2020.

[9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

[10] CCSDS Secretariat. CCSDS 120.1-G-3 Image Data Compression. Technical report, 2021.

[11] Diego Romano, Valeria Mele, and Marco Lapegna. The Challenge of Onboard SAR Processing: A GPU Opportunity. In Valeria V. Krzhizhanovskaya, Gábor Závodszky, Michael H. Lees, Jack J. Dongarra, Peter M. A. Sloot, Sérgio Brissos, and João Teixeira, editors, *Computational Science – ICCS 2020*, pages 46–59, Cham, 2020. Springer International Publishing.

[12] ESA. Copernicus Open Access Hub, 2023. Available at https://scihub.copernicus.eu/.

[13] Rich Hall. GitHub - Rich-Hall/sentinel1decoder: Python decoder for Sentinel-1 level0 files, 2022. Available at https://github.com/Rich-Hall/sentinel1decoder/.

[14] Georgios Pilikos, Mario Azcueta, Roberto Camarero, and Nicolas Floury. Raw Data Compression For Synthetic Aperture Radar Using Deep Learning. In *8th International Workshop on On-board Payload Data Compression*, Athens, Greece, 2022. ESA.

[15] Chané Pieterse, Warren P. Plessis, and Richard W. Focke. Metrics to evaluate compression algorithms for raw SAR data. *IET*

Radar, Sonar & Navigation, 13(3):333–346, March 2019.

[16] I.H. McLeod, I. G. Cumming, and M.S. Seymour. ENVISAT ASAR data reduction: impact on SAR interferometry. *IEEE Transactions on Geoscience and Remote Sensing*, 36(2):589–602, March 1998.

[17] F J Anscombe. Graphs in Statistical Analysis. *The American Statistician*, 27(1):17–21, 1973.

[18] Kwok and Johnson. Block Adaptive Quantization of Magellan SAR Data. *IEEE Transactions on Geoscience and Remote Sensing*, 27(4):375–383, July 1989.

[19] Joel Max. Quantizing for Minimum Distortion. *IEEE Transactions on Information Theory*, 6(1):7–12, 1960.

[20] Evert Attema, Ciro Cafforio, Michael Gottwald, Pietro Guccione, Andrea Monti Guarnieri, Fabio Rocca, and Paul Snoeij. Flexible Dynamic Block Adaptive Quantization for Sentinel-1 SAR Missions. *IEEE Geoscience and Remote Sensing Letters*, 7(4):766–770, October 2010.

[21] G. E. Hinton and R. R. Salakhutdinov. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313(5786):504–507, July 2006.

[22] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural Discrete Representation Learning. In *Proceedings of the 31st International Conferene on Neural Information Processing Systems*, pages 6309–6318, December 2017.

[23] Richard Hawkins, Colin Paterson, Chiara Picardi, Yan Jia, Radu Calinescu, and Ibrahim Habli. Guidance on the Assurance of Machine Learning in Autonomous Systems (AMLAS). Technical report, Assuring Autonomy International Programme, 2021.

[24] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end Optimized Image Compression. In *5th International Conference on Learning Representations*, 2017.

[25] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Conditional Probability Models for Deep Image Compression. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4394–4402, Salt Lake City, UT, June 2018. IEEE.

[26] David Steenari, Kyra Förster, Derek O'Callaghan, Maris Tali, Craig Hay, Mikulas Cebecauer, Murray Ireland, Sheila McBreen, and Roberto Camarero. Survey of High-Performance Processors and FPGAs for On-Board Processing and Machine Learning Applications. In *European Workshop on On-Board Data Processing*, June 2021.