SSC24-IV-09

Adaptive Raw SAR Data Compression Using Machine Learning Enhanced Block Adaptive Quantization

Craig Hay, Cameron Anderson, Lucy Donnell, Murray Ireland
Craft Prospect Ltd
Suite 12, Fairfield, 1048 Govan Road, Glasgow G51 4XS, UK; +44 7421 994 712
craig@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

Earth observation (EO) data plays a crucial role in climate monitoring, disaster response, asset management, and security, with synthetic aperture radar (SAR) data being particularly valuable. The ability of SAR to generate dense time series, unaffected by cloud cover or darkness, makes it ideal for monitoring applications. The commercial SAR sector has experienced significant recent growth, providing increasing amounts of data through advancements in quantity, quality, frequency, and dissemination. However, challenges in data management have emerged due to the volume of data captured by modern SAR instruments, surpassing satellite downlink capacities.

The Adaptive SAR Signal Compression Through Artificial Intelligence (SARAI) project, funded by the European Space Agency, adopted a unique approach to solving this on-board data bottleneck by focusing on enhancing the effectiveness of raw data compression using machine learning (ML) models. These models make inferences based on statistical features in the raw complex radar signals and select optimal compression algorithms based on content inferred from raw SAR data. The results are used to vary the encoding bitrate within algorithms and select between different algorithms, conserving bandwidth and improving system performance. This innovative strategy demonstrates the feasibility of extracting valuable insights directly from raw SAR data, a pioneering step in satellite data processing.

Traditionally, on-board processing of SAR data into images faces computational complexity challenges. This project's breakthrough lies in inferring information from data without creating SAR images and can be extended to applications such as target and change detection in the future. Implementing an ML model within the constrained, low-power computing environment of a satellite poses unique challenges. The model must be trained on representative data and operate effectively within these limitations.

The initial phase of the project involved a survey of publicly available raw SAR datasets, which informed the selection of relevant applications and frequency bandwidths for detailed study. A subsequent analysis and trade-off of current SAR compression algorithms involved evaluating these algorithms against various metrics such as applicability to specific bands, applications and complexity measures.

The project then assessed the feasibility and capabilities of machine learning in this context. A software prototype was developed to select algorithms and algorithm parameters based on statistical features of the raw data and compress data with different algorithms through the scene. Bit rate allocations were varied in Block Adaptive Quantization (BAQ) and FFT-BAQ algorithms. The results showed improvements over fixed bitrates, saving an average of 0.3 bits per block in the encoding, whilst maintaining similar signal-to-quantization-noise measurements in the image domain as the full BAQ 4-bit rate. Finally, there was an examination of hardware requirements and constraints. This was crucial in understanding the practical aspects of implementing ML-based compression algorithms in real-world SAR applications.

In conclusion, this project has laid a foundational framework for the future integration of machine learning in raw SAR data compression, offering a path towards more efficient and adaptive compression techniques that can significantly enhance the performance of SAR systems in various applications.

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 semi-conductor 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 Small- sat 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 stems. The work carried out in this project since the previous SmallSat conference in 2023 addresses this problem through the development of a novel on-board compression paradigm with results demonstrating the ability to infer image statistics from raw SAR data. In the following sections the raw SAR dataset collection and preparation survey will be outlined followed by the rationale for the standard compression algorithms that were selected. The machine learning (ML) solution will then be described along with the results that were obtained through the demonstrator. Finally, the hardware requirements and constraints for on-boarding this technology will be examined along with lessons learned and future work to be carried out in this domain.

RAW SAR DATASET PREPARATION

As the first stage of this work, a survey was conducted to collate available SAR mission data sources and find viable data or datasets to use in the development of on-board data compression algorithms.

Dataset Survey

Various historic and current SAR mission data was captured in a survey and analysed for suitability, with features, benefits and drawbacks documented. Table 1 outlines the availability of SAR mission data at the time of the survey. It is worth noting that since this survey work was concluded in March 2023, the availability of raw SAR data has increased with both Umbra and Capella starting Open Data Programs which include raw data in CPHD format and can be read with the *sarpy* tool.

The outcome of this was that Sentinel-1 data was selected within this activity for its free availability, worldwide coverage, good spatial and temporal resolution (with multi-polar data for some target areas) and thorough documentation. Acquisitions made in Strip Map Mode (SM) were specifically targeted due to the ease of image focusing compared to the more complex Interferometric Wave (IW) mode.

Backscattering and Signal Dynamics Analysis

A baseline source-coding algorithm for SAR raw data compression is the Flexible Dynamic Block Adaptive Quantization (FDBAQ), which is used in Sentinel-1. The main advantage of FDBAQ over BAQ is in using the adaptive quantization rate according to the local signal to noise ratio (SNR). To estimate the SNR, assuming

Table 1: Available mission data summary.

Mission	Raw (L0)	SLC (L1)	GRD (L2)	Raw Decoder	Global Coverage Raw Data	Multi Pol
Seasat	F	T	T	F	F	F
ERS-1	T	T	F	T	F	F
ERS-2	T	T	F	T	F	F
ENVISAT	T	T	F	T	F	T
Sentinel-1	T	T	T	T	T	T
RADARSAT-1	F	T	T	F	T	T
RADARSAT-2	F	T	T	F	T	T
CosmoSky MED	F	T	T	F	T	T
TerraSAR-X	F	T	T	F	T	T
Iceye	F	T	T	F	F	F
Capella	F	T	T	F	F	F
Umbra	F	F	F	F	F	F
SSTL NovaSAR	F	T	T	F	F	F

circular complex Gaussian clutter and thermal noise, *backscatter coefficients* of the raw data blocks must be calculated[4]. To conduct this task for comparison of current quantization methods, and later using for the AI based SAR raw data compression as a conditional input, we can either use models which incorporate electromagnetic (EM) scattering models, e.g. Born approximation EM model, or real SAR backscatters data[5]. We selected the latter approach here for a fair comparison. The Sentinel-1 Strip Map Mode I/Q data, i.e. before range or azimuth focusing, was used to estimate statistics of signal, i.e. backscatters.

The signal dynamics can be summarized as being affected by a combination of instrument properties, frequency band and by environmental factors

Environmental and Temporal Influences

1. **Terrain Types**: Impact varies across urban, agricultural, forest, mountainous, and desert areas.
2. **Weather Conditions**: Differences in signal dynamics under cloudy versus clear skies.
3. **Seasonal Variations**: Changes in signal characteristics across summer, winter, autumn, and spring.
4. **Diurnal Temperature Effects**: Variability in signal due to day and night temperature differences.
5. **Forest Types**: Distinct signal dynamics in rainforests compared to temperate forests.
6. **Urban Density**: Variation between urban and suburban environments.
7. **Building Materials**: Differences in signal reflection or absorption in steel/glass versus brick constructions.
8. **Temporal Changes**: Signal variations in the same scene captured at various times

Instrumentation and Methodology

1. **Polarization Effects**: Exploring signal differences in the same scene using various polarizations (VV, VH, HH, HV).
2. **Temporal Captures**: Repeated imaging of the same scene over time to observe temporal dynamics.

Approach for Capturing Signal Dynamics

1. **Time-Period Sampling**: Capturing the same scene across different time periods to assess temporal changes.
2. **Polarization Utilization**: Leveraging different polarization modes (VV/VH/HH/HV) to enrich data analysis.

Over a scene, the raw signal dynamics vary as can be seen with a selection of moments in **Error! Reference source not found.**. These can therefore be used to learn an adaptive coding scheme with different compression algorithms being used for distinct parts of a scene given the raw data statistics.

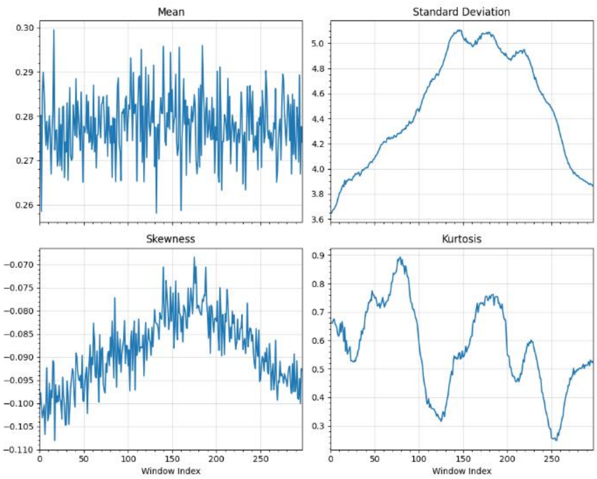

Figure 1: Signal moments over a scene.

To assess the suitability of data/dataset, a set of metrics were developed to both assess the SAR signal dynamics, and to assess the performance of developed compression algorithms. It is proposed to use an approach for assessing performance of compression algorithms similar to that found in[6].

Statistical Metrics – Data Domain

Standard statistical measures were calculated in the data domain. These are simple to calculate and provide some measure of how the data is changing pre and post compression. Although the raw SAR compression algorithms implement lossy compression, 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[7]. These models consist of the following set of assumptions:

- Complex IQ signal data, digitally sampled at e.g. 8-bits/sample
- Zero-mean, circular Gaussian distribution of complex samples
- 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 the 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 magnitude (I) and phase (Q) separately. The magnitude of the IQ signals therefore follows a Rayleigh distribution.

The metrics used to assess performance in the data domain were:

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

As the values of the raw data (and SLC focused images) are complex valued IQ signals, the polar form (magnitude and phase) were used to calculate these statistics.

Image Quality Measures

As well as signal statistics, image quality metrics can give an interpretable quantitative measure in the image domain. A non-exhaustive list of potential SAR image quality metrics is:

- Dynamic range – ratio of brightest vs darkest pixel
- Impulse response width (IRW) in both range and azimuth
- Peak-to-side lobe ratio (PSLR) in range and azimuth
- Image contrast – the ratio of average intensity of distributed clutter background (standard deviation of image) over the average intensity of no return background (mean of image).
- Global contrast factor – A weighted sum of local contrasts of a range of smaller image sizes
- Mean Squared Error – Calculate for difference in amplitude of the original image and post-compression image
- Mean Phase Error – Calculate for difference in phase of the original image and post-compression image
- Signal-to-distortion Noise Ratio – Calculate for difference in amplitude of the original image and post-compression image

SELECTION AND TRADE-OFF OF SAR COMPRESSION ALGORITHMS

An investigation and analysis of traditional algorithms used for SAR data compression was conducted for use in the selective compression algorithm. Following the implementation of selected algorithms, a pipeline was developed to create a machine learning ready dataset by processing raw SAR scenes with the various compression algorithms and deriving input features and target variables.

Algorithms Survey

The standard algorithms for SAR data compression were selected for comparison, these are block adaptive quantization (BAQ) and its variants: Entropy Constrained BAQ, Block Adaptive Vector Quantization, Flexible BAQ, etc. These improve BAQ but add complexity. The survey found most algorithms are based on an underlying statistical model and set of assumptions for SAR data:

- **Characteristics**:
  - 8-bits/sample complex data.
  - Zero-mean, circular Gaussian distribution of complex samples.
  - Minimal saturation in ADC.
  - Low correlation in I and Q channels and intersample.
  - Variance changes slowly in slant range and azimuth.

- **Model suitability**: Particularly effective for homogeneous scattering regions.
- **Data representation**: Complex valued (I and Q), can be converted to polar form (magnitude/phase).
- **Impact of compression**
  - **Objective**: Lossy compression should not significantly alter overall data statistics.
  - **Statistical measures**: Standard measures will be calculated in the data domain.
- **Conclusion**
  - **Simplicity and efficiency**: BAQ serves as a simple yet effective base algorithm.
  - **Future work**: Exploration and categorization of BAQ variants for further optimization.

The relationship between BAQ and its descendants/variants is illustrated in Figure 2 below.

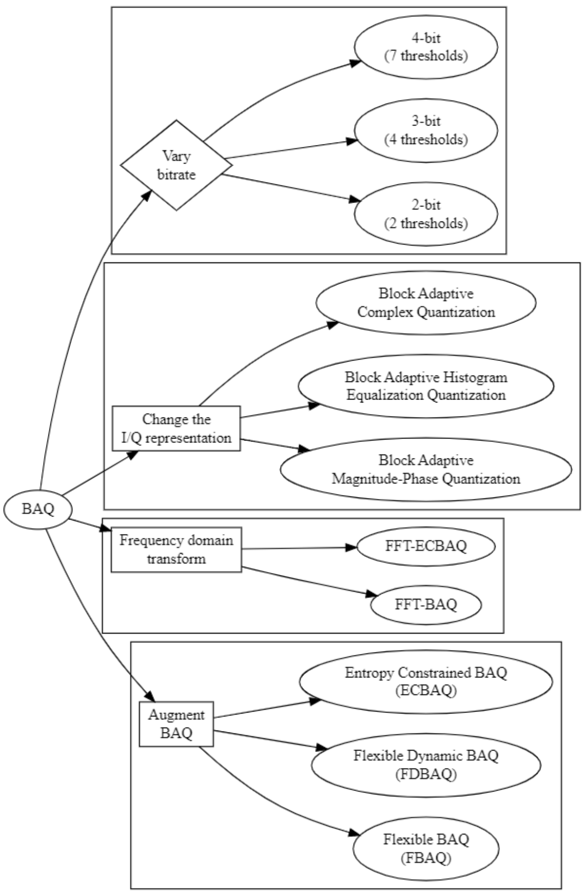

Figure 2: Algorithm taxonomy.

The benefits and limitations of these algorithms were analysed. Trade-offs between algorithms can be made on grounds of implementation complexity, memory usage, CPU usage, hardware implementation complexity, patent issues. Qualitative trade-offs are outlined below in Table 2.

Table 2: Qualitative trade-offs.

Algorithm	Pros	Cons
BAQ	Simple to implement in hardware Well tested and understood	Must send side information (block variances) which reduces the compression ratio Must choose a block size when designing implementation Coding delay -- must see whole block before quantizing Integer rates mean adjustments can only be made in 3dB steps in SQNR
EC-BAQ	Higher compression ratio than BAQ Non-integer code rates possible Extended instantaneous dynamic range Backward adaptive – quantization parameters are chosen based upon the history of past samples	Increased complexity of entropy coder
FFT-BAQ	Further compression gains of 20% on BAQ by zeroing low frequency coefficients Hardware accelerated FFT implementations on modern hardware	High computational load to perform FFT onboard
FFT-ECBAQ	Further compression gains of 20% on ECBAQ by zeroing low frequency coefficients Hardware accelerated FFT implementations on modern hardware	High computational load to perform FFT onboard

Following the survey, the following algorithms were then implemented and analysed:

- BAQ
- FBAQ
- FFT-BAQ

Data Labelling

A machine learning ready dataset was created by compiling a set of representative data from ESA Scihub and generating features and target variables for training the machine learning model. These target variables were generated by applying various implemented SAR raw data compression algorithms to the raw data and focusing into images. The original and reconstructed images were then used to generate the target variable, which was going to be an error metric between the original and reconstructed data in the image domain.

The steps taken to generate the training data are listed below

1. Search scenes using Browser and filtering on Copernicus Data Space.
2. Get Zipper URLs for selected scenes.
3. Script downloads through API.
4. Decode raw products using *sentinel1_level0* decoder.
5. Compress/decompress original raw product for both polarizations using Python implementations of raw SAR compression algorithms. Save as .npy files.
6. Focus original products and those reconstructed with each implemented compression algorithm.
7. Split scenes into decision windows (selected for this activity to comprise of forty range lines).
8. Generate machine learning features from original raw products and save as CSV.
9. Generate target variables by calculating error metrics in image domain between original raw data and each set of reconstructed raw data. Save as CSV.
10. Training data generation complete – CSV of features and target variables.

The methodology used to generate features and metrics is illustrated in Figure below. The window size carried out in step 7 above was selected ad hoc to ensure balance between an unnecessary amount of switching between compression algorithms – which would require more data to be sent containing the compression algorithms used – and having enough variations between the windows to allow for beneficial algorithm selection. This is a parameter that requires further investigation to optimise the window size for training.

he different training data that was used was made up of available SM targets. There is a limited selection of SM targets available, as it is only used operationally to capture small island targets where the IW swath is far too wide.

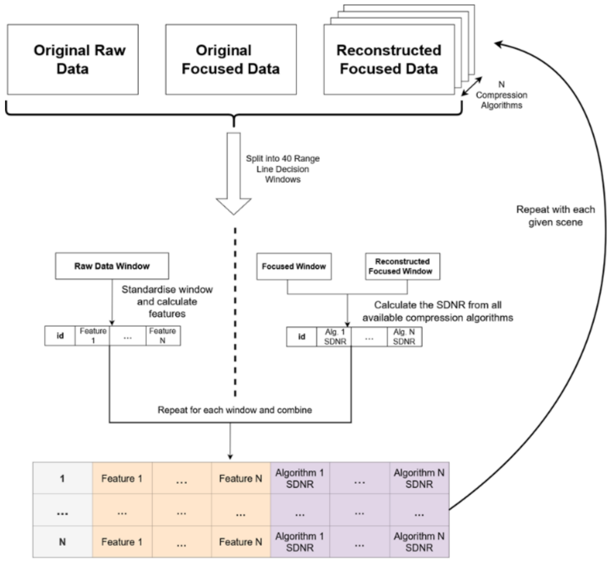

Figure 3: ML-ready dataset construction methodology.

AI/ML FEASIBILITY AND CAPABILITY ASSESSMENT

After having carried out dataset curation and gaining and understanding of traditional algorithms used for SAR data compression, the feasibility of utilising ML for SAR data compression was explored. An initial approach was outlined to experiment with VAE and LSTM models and later revised to an approach that utilised ML to infer the best traditional compression algorithm choice based on the contents of the scene. This was defined in a problem statement as:

Raw SAR data has a slowly changing variance in both range and azimuth for the I and Q signals. Current raw SAR compression algorithms compress the raw data blockwise, with blocks being small enough that it can be assumed that the variance is not changing within a block e.g. 128 consecutive values in range or combination of range and azimuth.

The algorithms used for compression are optimal when blocks exhibit Gaussian statistics. Due to the nature of SAR backscatter this assumption holds closest to the truth when the scene is homogeneous.

These statistics are caused by the constructive and destructive interference in reflections from random scattering in each range cell. These Gaussian assumptions do not hold when the scene is inhomogeneous e.g. in the presence of a point scatterer.

Therefore, the ability to switch between compression algorithms based upon the contents of the scene is seen as advantageous. These scene contents can be inferred from the statistics of the backscattered signal. Machine learning inference can be employed to predict metrics *that can be mapped to the presence of "cultural clutter" in the scene.*

Feasibility Assessment

This work considered the feasibility of the ML model to improve on traditional methods of onboard data compression, considering the ML model and input feature complexity, metrics for data quality and the onboard data processing pipeline.

Simple statistical features were selected as candidate features for the ML model. Through testing, it was discovered that these simple features allowed for a well performing model for predicting the error in the reconstructed image due to the chosen raw data compression method. Many of the features chosen would already be calculated for the standard onboard compression algorithms, so it is feasible that these features or similar can be generated from the raw data on-the-fly at the rate of data acquisition (if necessary). Features investigated were:

- Arithmetic Mean ($\bar{y}$)
- Standard Deviation (σ)
- Skewness – A measure of the asymmetry of the distribution of the given data section
- Kurtosis – A measure of the data contained within the tail of the distribution as compared to a normal distribution
- Normal Test – A combination of the z-scores of a skewness (s) and kurtosis (k) test of the form $s^2 + k^2$ which tests the null hypothesis that samples come from a normal distribution
- Kolmogorov-Smirnov Test – A test of the equality of two distributions which tests how likely it is that the data sample is drawn from a normal distribution
- Shapiro-Wilk Test – Another test statistic to test the null hypothesis that the data sample is drawn from a normal distribution

Through a process of recursive feature selection, a final set of features was reduced to:

1. Kolmogorov-Smirnov test, I channel, cross polarization
2. Kolmogorov-Smirnov test, Q channel, cross polarization
3. Kurtosis, I channel cross polarization
4. Kurtosis, Q channel co-polarization
5. Normal test, Q channel, Cross polarization
6. Skewness, Q channel, co-polarization

This number of features was selected as it maximized model accuracy after evaluating the best features for the whole range of feature numbers, as shown in Figure 4. It was also observed that model accuracy did not increase with the increasing number of features so it is preferable to select a smaller number of features for model simplicity and to reduce computational complexity when performing on board inference, as these features would have to be generated from the IQ signals on-board.

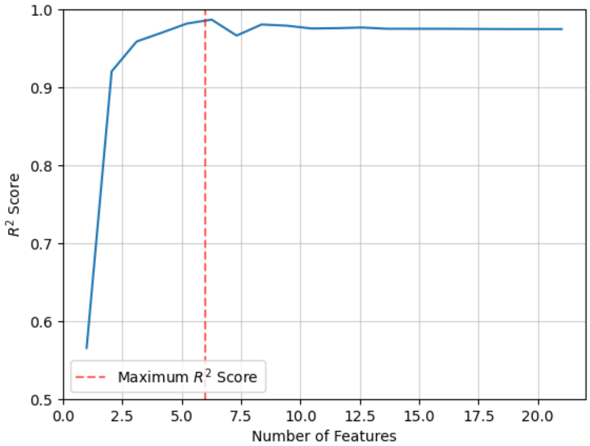

Figure 4: R2 score against number of features showing the optimal feature number.

Notes on the machine learning model implementation are shown below.

- **Application Scope**: The algorithm selector is not specialized for applications like polarimetry or interferometry.
- **Final Output**: Aimed at optimizing the Synthetic Aperture Radar (SAR) Single Look Complex (SLC) image.

Target Variables for Image Evaluation

- **Initially Considered**: Phase Error, Mean Squared Error (MSE), Error dB, Signal-to-Distortion Ratio (SDNR).
- **Selected Metrics**: MSE and SDNR were chosen for their general applicability.

$$MSE = \frac{1}{MN} \Sigma_{x=1}^{M} \Sigma_{y=1}^{N} \left(g(x,y) - f(x,y)\right)^2 \quad (1)$$

Where $g(x,y)$ and $f(x,y)$ are the original and reconstructed $M \times N$ images respectively.

$$SDNR = 10 \log_{10} \left[\frac{\Sigma_{x=1}^{M} \Sigma_{y=1}^{N} g(x,y)^2}{\Sigma_{x=1}^{M} \Sigma_{y=1}^{N} \left(g(x,y) - f(x,y)\right)^2} \right] \quad (2)$$

Table 3: ML models considered and their preliminary results.

Model	Pros	Cons	Testing Hyperparameters	R^2
K-Nearest Neighbours	• Inherently multi-output • High predictive power • Easily interpretable output • Well suited to regression problems	• Requires the storage of training data for predictions • Can be computationally expensive and slow to predict for large datasets	• N-neighbours = 5 • Weights = uniform • Algorithm = Ball Tree with leaf size of 30 • Distance metric = Euclidian	0.864
Decision Tree	• Inherently multi-output • Less data preparation required – although in this case, data is required to be standardized to compare statistics • Can be used to identify the most important features and down select • Easily interpretable output	• Prone to overfitting with reduction of training set error at the expense of test set error – solved through pruning and constraining parameters of the model • Can become complex to calculate predictions as tree grows leading to longer training times also	• Split quality metric = Squared error • Max depth = None • Minimum samples to split a node = 2 • Minimum samples to be a leaf = 1 • Maximum features to look for best split = None • Maximum leaf nodes = None	0.801
Support Vector Regression	• Scales well to high dimensional data • Less risk of over-fitting compared to other models • Can solve complex problems with appropriate kernel function selection	• Not inherently multi-output so requires a wrapper so handle this • Longer training time compared to the other algorithms • Difficult to choose appropriate kernel	• Kernel = Radial Basis Function • Gamma = 1/(number of features * variance in features) • C = 1 • Epsilon = 0.1	0.268
Linear Regression	• Inherently multi-output • Simple implementation and short training time • Not computationally complex in predictions	• Assumes a linear relationship between independent and dependent variables so cannot handle more complex relationships • Outliers have a significant impact on the regression		0.729

Rationale for Chosen Metric

- **Advantage of SDNR Over MSE**:
  - **MSE Limitations**: Does not specify if errors are many and small or few and large; strongly influenced by intensity scaling, making it unreliable for comparing different scenes.
  - **Impact on ML Model**: The reliance on standardized data in ML reduces MSE's effectiveness as a predictive metric.
  - **SDNR Benefits**: Offers a more comprehensive and reliable measure, overcoming MSE's shortcomings; normalises the MSE to form a more global measure of error.

- **Metric Selection Conclusion**: SDNR is preferred for its ability to provide a more accurate and consistent measure of image quality across various scenes, making it a more suitable target for the algorithm selector in SLC image processing.

- **Purpose**: Predict the "expected error" (SDNR) for various compression algorithms upon analysing raw data.
- **User and System Requirements**: Incorporate user-defined factors and system tolerances.

Machine Learning Dataset Creation

- **Image Sets**: Included raw original data, focused original, and data compressed using BAQ2, BAQ3, BAQ4, FFTBAQ2, FFTBAQ3.
- **Segmentation**: Data split into forty range line decision windows for applying a single compression algorithm.

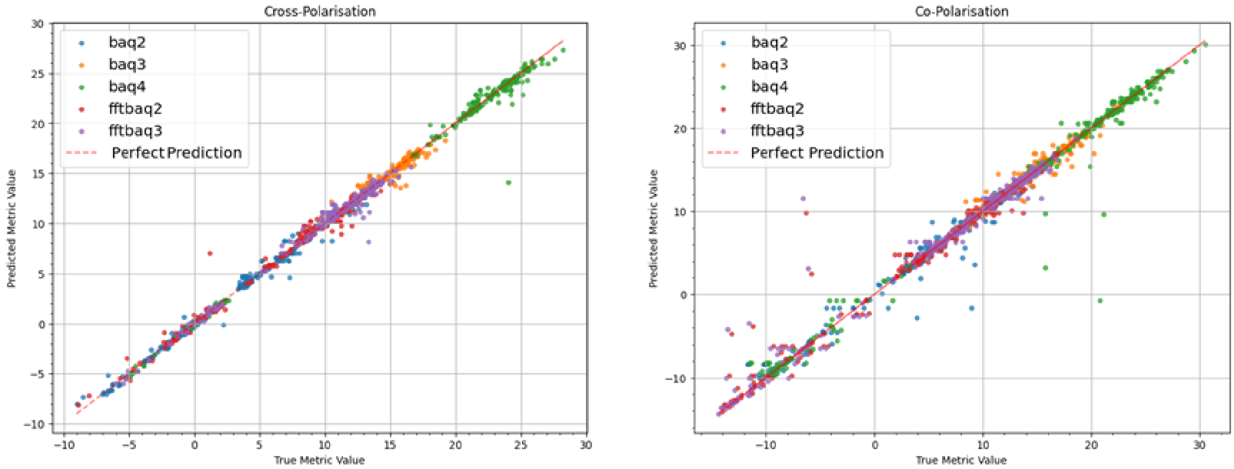

Figure 5: Predictions against the true metric values for all compression methods.

- **Feature Calculation**:
  - From raw data: Features described in Section **Error! Reference source not found.** calculated for regions of forty range lines.
  - From focused images: SDNR values calculated for corresponding forty range line windows.
- **Dataset Composition**: Combined features and SDNR values form the full dataset for ML predictor model training.

Conclusion

- **ML Model Objective**: To train a model that can predict the SDNR for different compression algorithms based on raw SAR data features.
- **Importance of SDNR**: Chosen as the primary metric due to its robustness in evaluating the quality of compressed images across varied scenes.

Several models were then assessed for implementation. Trade-offs for the models that were under consideration are shown in Table 3**Error! Reference source not found.**. This shows the pros and cons of the different models alongside an initial accuracy score chosen to be the coefficient of determination (R^2).

Selecting the two best initial scoring models (decision tree and k-nearest neighbours), initial model results are shown in the next subsections. Also detailed are the hyperparameter grid searches that were carried out for these models, as well as feature down-selection using recursive feature elimination.

Decision Tree Results

Described here are the results from the Decision Tree predictor method. Figure 5 displays the predicted values against the true values that were previously calculated as part of the test dataset. A dotted line shows perfect predictions. Accuracy of the predictions is 0.974 in the R^2 score and 1.09 in mean squared error over the test dataset.

A grid search was carried out to avoid model overfitting with a parameter grid focusing on tree depth and how the tree is split. The final parameters selected were:

- Maximum tree depth = 15
- Minimum number of samples required to be left at a leaf node = 2
- Minimum samples required to split an internal node is five.

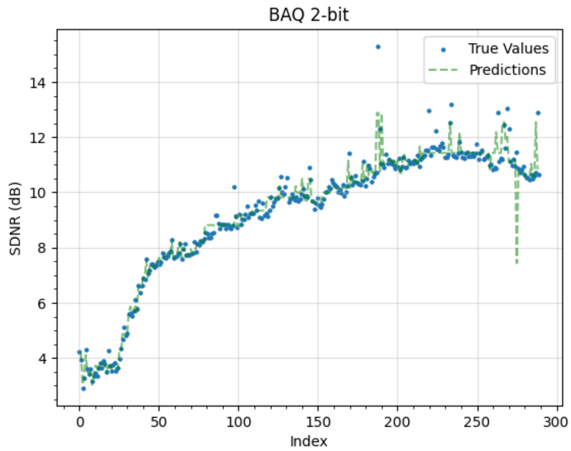

Figure 6: Predictions from decision tree algorithm over a scene for lowest bit rate BAQ.

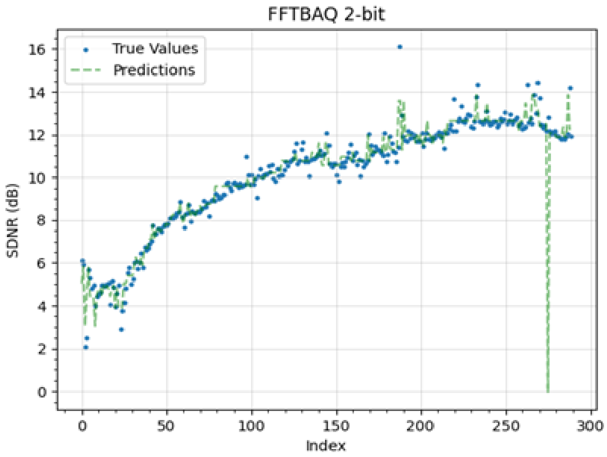

Figure 7: Predictions from decision tree algorithm over a scene for lowest bit rate FFTBAQ.

The above figures show scatter plots of the true SDNR values with the decision tree predictions plotted on top. These are taken over a full scene which in this case is from Houston showing the SDNR for the co-polarisation. The spikes seen in the data are mis-predictions from the machine learning model where the model may still be overfitting to the training data. Although these may cause single range windows to be compressed with a non-optimal compression algorithm, the model accuracy over the whole scene will ensure that the correct compression algorithm will be chosen most of the time thus incurring the benefit of using the ML based approach. These predictions could be improved by increasing the amount of training data with a wide range of clutter in the scenes allowing the entire range of input space to be seen by the model. It is also yet to be investigated if the size of the range line window would affect the accuracy of the model, as there could be range lines in these windows which are dominating the average used for the feature vector and thus causing misdetections.

Decision Module Design

To be able to select the correct compression algorithm, a post-processing step was implemented to combine the output of the ML model and user defined parameters such as the error threshold, the desired bit rate and computational complexity weightings. This was done with the aim of achieve a quantization error that is equal to or less than a user defined threshold while minimising the bitrate or computational complexity of the compression. The effects of compression are negligible if the SDNR is equal to 25dB, so this is defined as the maximum SDNR value corresponding to 'perfect' compression. Bitrate and computational complexity weightings are used to set an order of preference of algorithms to be selected, depending on the given use case. To achieve the desired bitrate, logic was implemented to simply choose the lowest bitrate algorithm that meets the SQNR threshold. In this way, the bitrate is kept to a minimum within the integer bitrate choices available.

Performance Assessment

Testing was conducted to assess the compression performance gains when using the ML enabled decision module. Results from the ML solution are presented below and compared with those gained when using only traditional algorithms. For this test, an Aurora scene was used at an unseen time with a different ground track from the training data, also using HH/HV polarisation.

Test Scene 1 – Aurora

For the initial test, an Aurora scene was used at an unseen time with a different ground track from the training data, also using HH/HV polarisation as opposed to the training data which was a VV/VH acquisition.

Scenario 1 – error threshold

The aim of this test was to show that the predictor and decision function could switch algorithm whilst maintaining an overall SDNR figure above the set threshold. It would also give an indication of where and how switching would occur, with the initial hypothesis that the algorithm would switch to lower bitrates in area with:

- Error threshold = 15dB
- Complexity weight = 0
- Bit-rate weight = 0

The average bit rate and corresponding SDNR values are shown in Table 4 with the decision boundaries shown overlaid on the scene in Figure 8. In this Figure, it is seen that, over the cultural clutter of a city where there are more backscattering present, the machine learning approach selects the BAQ4 algorithm to maintain the accuracy of the reconstructed image. This then changes to the BAQ3 algorithm over the fields where the same accuracy can be maintained while using less bits to encode the data.

Scenario 2 – weighted bitrate

Using this same Aurora scene, the scenario of a constrained bitrate was tested. This scenario was designed to demonstrate the decision function meeting multiple constraints, with a minimum error threshold to be maintained whilst keeping the average bitrate below a certain level. This was designed to exercise the use of more computationally complex algorithms which could meet these dual constrains of minimising both error and bitrate whilst trading off on increased computation. The decision function parameters were set to:

Table 4: SDNR and bitrate results comparing single compression algorithms to the ML-based approach.

	Cross Polarization		Co-Polarization	
	SDNR	**Bitrate**	**SDNR**	**Bitrate**
ML Approach	24.24	3.68	20.22	3.69
BAQ3	13.88	3	16.3	3
BAQ4	25.10	4	24.44	4

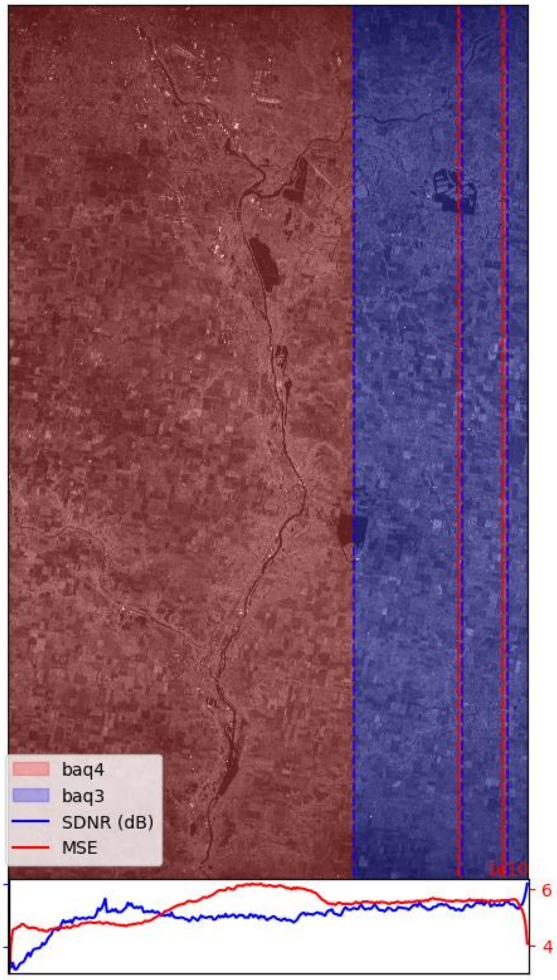

Figure 8: Scene showing which compression algorithms are selected over different regions. It is seen that over regions with cultural clutter, the BAQ4 algorithm is used whereas when there are only fields the less accurate BAQ3 is used to achieve the same SDNR.

The decision function parameters were set to:

- Error threshold = 15dB
- Complexity weight = 0
- Bit-rate weight = 4.5

This indeed causes the decision module to penalise higher bit-rate algorithms in favour of the lower bitrates. This also highlights the use of the FFTBAQ algorithms being used as they can achieve a higher accuracy whilst maintaining a lower bitrate, although this comes at the cost of higher computational complexity. Where FFTBAQ was used here, it was predicted to still meet the set error threshold of 15dB whereas BAQ2 was not.

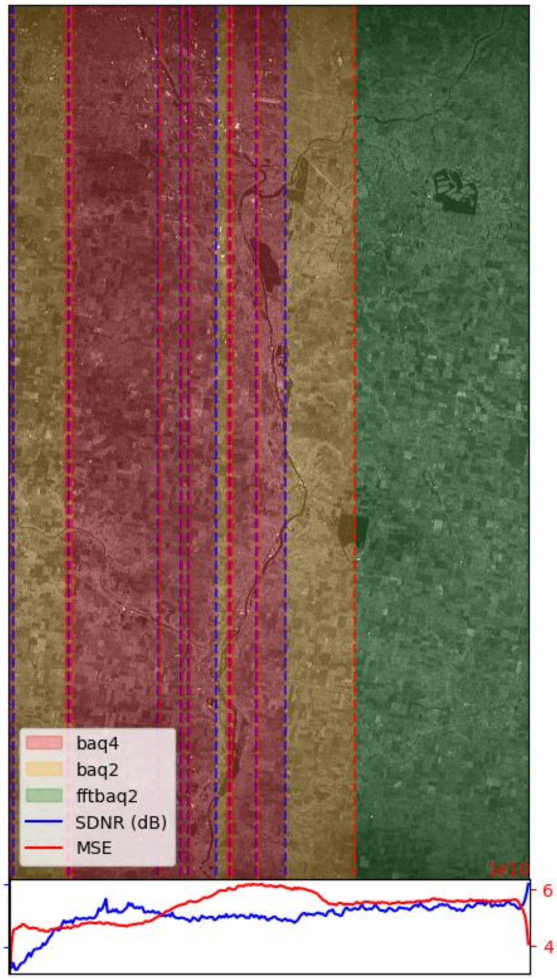

Figure 9: Scene showing the compression algorithm selections when the decision was weighted towards lower bitrate use. The higher bitrate is allocated to the regions with most cultural clutter.

Test Scene 3 – Saint Helena

This island scene was selected to test as an unseen region which acts as a simple and clear demonstration of the selection method.

The parameters used for this test were:

- Error threshold = 18dB
- Complexity weight = 0
- Bit-rate weight = 0

Since this scene contains a large amount of ocean with an island in the middle, it would be preferential to have a compression scheme which compresses the ocean with less bits and is able to detect the island when it occurs and compress this with a larger number of bits, which is what is observed in the Figure below. It is seen though that the selection of BAQ4 at the right-hand side of this scene is due to ocean current that can be seen in the VV and VH dB image causing greater backscattering and thus a higher bit rate selection.

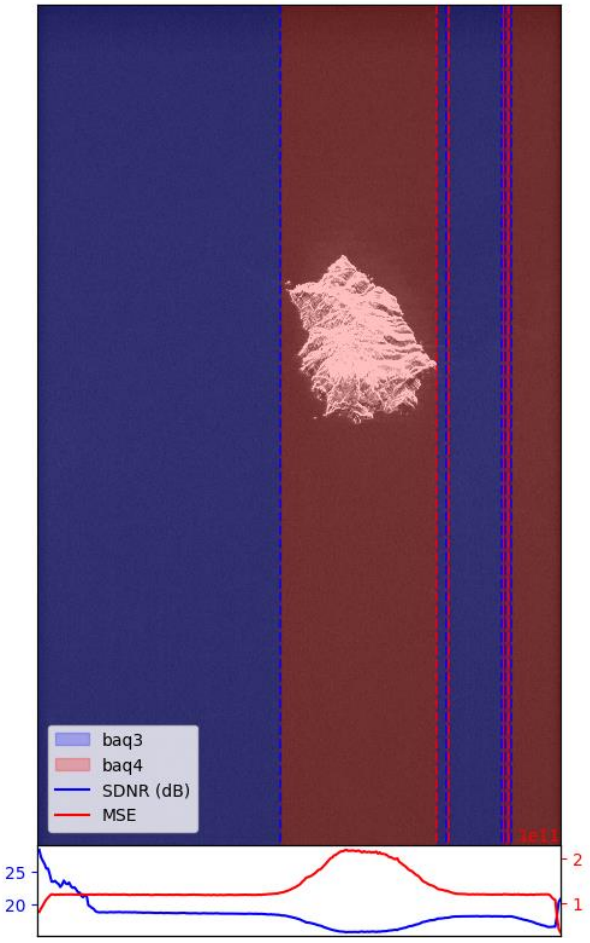

Figure 10: Island scene again showing the higher bitrate being selected for the island to maximise the accuracy whilst minimising the bitrate for data transmission.

HARDWARE FEASIBILITY

This review explored advancements in on-board processing for small satellite (SmallSat) scale SAR missions, focusing on the emerging NewSpace sector. The trend in these SAR satellite systems is shifting towards deploying large constellations, enhancing robustness, reliability, and performance, including faster tasking and data delivery. These constellations allow for immediate analysis and re-tasking of subsequent satellites for targeted imaging. SAR platforms, particularly SmallSats, face challenges due to heavy duty cycling, with limited duty cycles significantly impacting operation. Synspective's data[8] indicates typical duty cycles of 5-10% per orbit. Increasing satellite numbers can alleviate this, but also imposes higher demands on ground station infrastructure. The active radar sensors in SAR platforms require substantial power for imaging, leading to significant heat generation and thermal management challenges on these resource-limited platforms. Efficient data handling is crucial, given the high data acquisition rates. Enhancing on-board data compression can address these challenges by reducing storage and downlink needs. Potential advancements include prioritizing data for downlink, satellite tasking based on extracted information, and change detection, leveraging advancements in on-board data processing and compression capabilities.

Baseline Hardware Architecture

In this section, a short case study for a potential future target system is presented. The target system is NovaSAR-S, which is an S-band SAR SmallSat owned and operated by Surrey Satellite Technology Limited (SSTL). NovaSAR provides a suitable candidate for comparing a runtime compression mode with an offline compression operating mode, as it currently operates runtime compression and SSTL are actively exploring offline on-board payload data processing architectures. One such project is the InCubed funded Flexible Intelligent Payload Chain (FIPC) where SSTL partnered with Craft Prospect Ltd and University of Surrey as application developers for an offline on-board processor designed to integrate the platform to process data from optical EO payloads.

The NovaSAR SAR instrument is based on the Airbus DS New Instrument Architecture (NIA)[9]. Regarding its physical attributes, the NIA Back-End unit weighs 11kg and consumes 50W of DC power. Inside of the NIA Central Electronics we are most interested in the Digitiser module as this is where the compression functionality currently sits. In fact, this is where all digitization and packetization functions are performed. The Digitiser samples the RF signals with a 12-bit ADC and encompasses all necessary digital signal processing and data handling for the receive chain. These steps performed on the digitised data are:

1. Digital down-conversion
2. Flexible L/M decimation filtering
3. Data compression options (2.5, 3, 3.5, 4, 4.5, and 5-bit Block Adaptive Quantisation)
4. Packetisation as per CCSDS standards

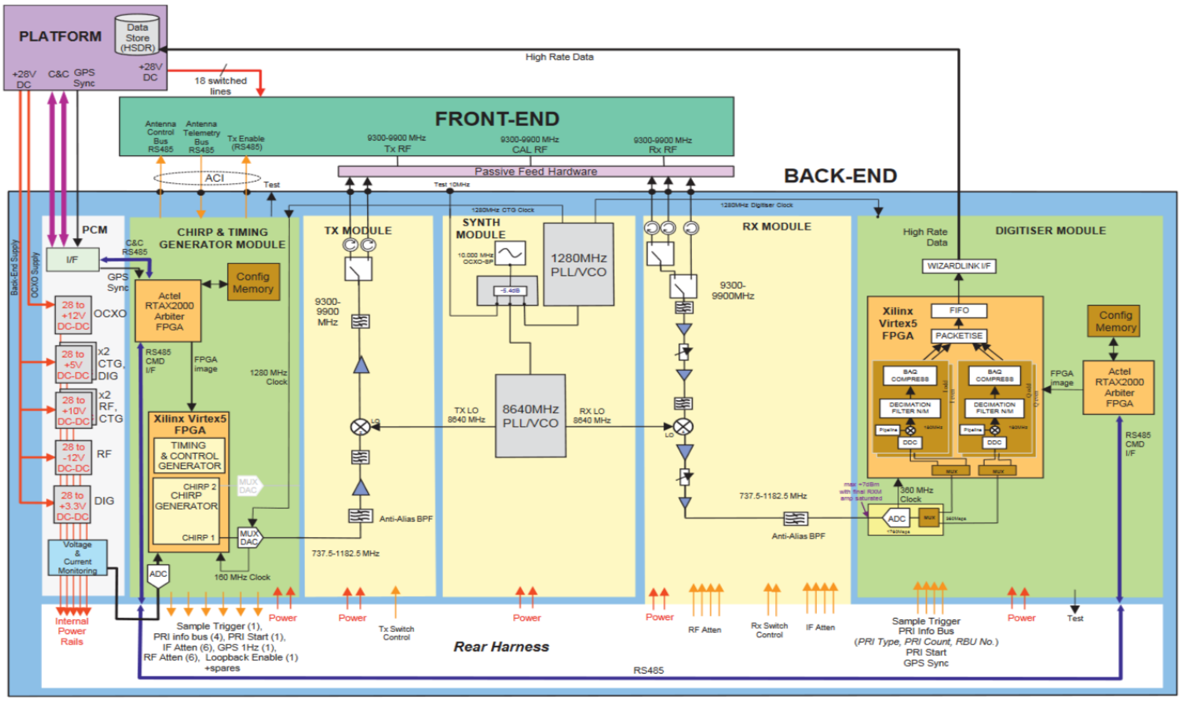

Figure 11: NIASAR-X payload NIA central electronics block diagram[10].

5. Datation features including time stamp, PRI count, packet count, Mode ID, PRI ID, etc.
6. Buffering and output to a high-rate data interface

An architecture diagram is shown in Figure 11. The details of the Digitiser FPGA can be viewed in larger scale in Figure 12. This module implements most of the required functionality on the Xilinx Virtex 5 (XQR5V) FPGA. The NIA also includes a highly flexible Chirp & Timing Generator (CTG) Module incorporating another Virtex 5.

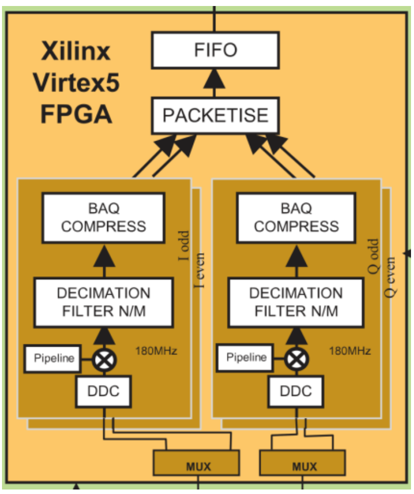

Figure 12: Digitiser architecture.

It can be seen from the Digitiser architecture that the I and Q signal data has been multiplexed into I even, I odd, Q even, Q odd streams of samples. This suggests that the decimation and BAQ compressors have been implemented in parallel to increase throughput.

The referenced thesis[11] demonstrates the efficient implementation of the Block Adaptive Quantization (BAQ) algorithm on FPGA hardware, leveraging MATLAB/Simulink, Xilinx System Generator, and Xilinx ISE. This approach, utilizing a Xilinx Virtex-4 FPGA, proved more efficient than a full VHDL implementation, enabling rapid development and verification of hardware prototypes for on-board compression. The FPGA's low resource utilization, using only 3% of available logic slices and 34% of block RAMs (mostly for test data), indicates substantial capacity for more complex or multiple data compression algorithms. This efficiency is further evidenced in the NIA Digitiser case study, showing that most digital functions of a radar backend can be integrated onto a single chip. This FPGA-based approach is promising for real-time SAR data processing, allowing for advanced processing capabilities with efficient hardware resource management.

Runtime compression operating mode

To implement a runtime operating mode, an estimate of the data to be processed is required to estimate the required throughput of any additional processing. This is a function of pulse repetition frequency (PRF) and range line length. The PRF is mission and system design dependent and varies based on satellite mode. A higher PRF typically corresponds to a narrower swath and greater azimuth resolution and vice-versa.

For Sentinel-1 SM, the PRF in SM is approximately 1500Hz. This results in 1500 range lines, each consisting of thousands of IQ pairs, being generated per second at the radar backend. These are typically 8/10/12-bit pairs of values depending on the ADC and digitization strategy. For a range line with around twenty thousand range bins, as is the case in Sentinel-1 SM, this gives a requirement of:

Range lines x IQ x ADC bits x range line length
1500 x 2 x 10 x 20000 = 480Mbps

For NovaSAR-1 the PRF varies between 5000-7000Hz. For an example in ScanSAR mode (shown in 13) with PRF 5302Hz and range line length 1117 this gives a data rate:

5302 x 2 x 12 x 1117 = 142Mbps

These are both high bitrates to service in real-time

Operational Mode Name	20m_ScanSAR_25km_HHHV_N6
Acquisition ID	12716
Mode Mnemonic	45m Dual Pol SCD
Raw Data Start Time	2020-07-05 23:05:10.18773
Polarisations	HH HV
Number Of Swaths	1
Pulses Received Per Dwell S1	1117
Number Of Pulse Intervals Per Dwell S1	1140
Rank S1	23
Receive Gain S1 (dB)	-12.42
Radar Centre Frequency (Hz)	3.2E+09
Pulse Repetition Frequency S1 (Hz)	5302.23
Pulse Length S1 (sec)	4.6122E-05
Chirp Bandwidth S1 (Hz)	5.0280E+07

Figure 13: NovaSAR acquisition parameters[12].

The main issue with the runtime compression mode is that most SAR satellite system architectures have the onboard compression sitting somewhere in the radar backend. For instance, SSTL's NovaSAR-1 contains a Digitiser module that carries out analogue-to-digital conversion of the RF signal, through multiple processing functions, to packetization with CCSDS format packets ready for transfer through high-speed serial to "data-recorder" and onward to downlink. This means that to modify this system to include ML-driven on-board compression, it would be the Digitiser module that would have to be modified.

To run inference before compression, suitable buffering would have to be in use. It is likely that this buffering would encompass the previous, current and next set of range lines to be processed. One buffer would be filled by the active sampling of the ADC, another with the data to run inference on, and another with the data being worked on for compression. With the chosen inference window of forty range lines, this equates to:

N range lines x range line length x IQ bytes x N buffers

40 x 20000 x 2 x 3 = 4.8e6 bytes = 4.8MB

The current window to completely process the data held in a buffer would be, in the case of Sentinel-1 Strip Map mode with a PRF of around 1500Hz:

Window = range_lines_per_window / range_lines_per_second

Window = 40 / 1500 = 26.67ms

This kind of analysis would need to be repeated for the target mission and clearly a different range line window would be chosen for NovaSAR.

The main advantage of the runtime compression configuration is that no additional mass memory is required in the architecture to store SAR data that needs to be worked on by the inference engine. As most of the digitization functions are implemented on a single (albeit large) FPGA, the question remains whether there is adequate FPGA resource remaining to implement the feature generation, inference and algorithm selection steps, as well as multiple SAR compression algorithms.

Offline Operating Mode

Addressing the challenges of implementing an offline processing mode in SAR satellite systems involves significant architectural changes, particularly in data compression within the radar backend. Currently, these systems typically have a bypass mode that avoids data compression, but this approach requires substantial mass memory to store uncompressed data from the ADC. For example, the SSTL Digitiser not only compresses but also processes data up to packetization for downlink, necessitating correctly sized mass memory for real-time buffering. A modification could involve constantly running the compressor in bypass mode and directly packetizing data into mass memory, though this would demand significantly more storage space. Another innovative approach under consideration is using a compressed data format (like BAQ4) as the standard. Here, a machine learning system would read data out from mass memory and evaluate if the data can be further compressed without notable loss in image quality. This method could potentially eliminate the need for a bypass mode, as data would be decompressed and recompressed at a lower bitrate when feasible, thus reducing the mass memory requirements. However, this approach raises uncertainties regarding mass memory layout or fragmentation and the simplicity of depacketization, as data are packetized following CCSDS standards by the radar backend. These considerations highlight the complexity of integrating advanced data compression techniques into existing SAR satellite architectures.

The hardware architecture for a system supporting multiple compression algorithms, such as BAQ, presents an intriguing design challenge. For each BAQ bitrate, a corresponding lookup table of quantization thresholds must be stored, with the number of values ranging from three for 2-bit BAQ to fifteen for 4-bit BAQ. If standard deviation calculations and data normalization can be conducted on-board, these tables can remain compact. Alternatively, storing a lookup table for each possible standard deviation (within a margin of error) could eliminate the need for on-board block normalization, trading computation for storage—a classic space/time complexity decision that can be left to the hardware designer.

These tables could also be utilized for the FFT-BAQ quantization process, with modern FPGAs capable of accelerating FFT operations. A key consideration is whether each compression algorithm should be implemented on separate boards, which can be individually powered on and off. This decision impacts the benefits of the scheme, especially in terms of power savings and computational complexity. If all algorithms are implemented on the same FPGA, it becomes crucial to assess the power-saving benefits of the scheme. Preliminary research indicates that strategies like clock gating could significantly reduce the FPGA's dynamic power usage[13], though the overall impact on the power consumption of the Digitiser remains uncertain. This complexity underscores the need for careful design considerations in developing efficient and effective compression systems for satellite hardware.

When considering hardware for space applications, several critical factors must be assessed, including space environmental effects, device size, weight, power (SWaP) requirements, system-level integration, and data processing throughput. The choice between radiation-hardened by design (RHBD), radiation tolerant (RT), or commercial off-the-shelf (COTS) devices depends on these factors, along with the platform’s power budget, total mission cost, and risk tolerance. RT and COTS devices often provide better performance per watt compared to RHBD devices, but there is a trade-off between radiation hardness and capability, as devices with larger process nodes are less susceptible to radiation-induced Single Event Upsets (SEUs).

Newer COTS devices, which are smaller and more efficient, are increasingly being adapted for space environments. They offer significant performance improvements, particularly in accelerating machine learning workloads on edge devices, making them suitable for small satellite power budgets. Examples include the Xilinx Ultrascale+ MPSoC and Intel Myriad X. These advancements are complemented by device-specific tools like Intel OpenVINO and Xilinx Vitis AI, which facilitate the integration of trained models onto these platforms. The maturity and extensive user base of COTS development tools are also advantageous.

The choice of specific devices for future demonstrations will be influenced by the processing system architecture, derived from the reference data processing architecture. Devices with native fault-tolerance, such as RHBD with integrated error-checking, require less system-level fault tolerance. Conversely, less mature devices like the Myriad X VPU necessitate more robust system-level fault mitigation strategies, including advanced fault detection, isolation, recovery (FDIR) systems, and potentially component duplication or triplication. This selection process underscores the balance between leveraging advanced capabilities of newer devices and ensuring reliability and resilience in the challenging space environment.

Feasibility of Implementing Inference Functionality

The raw data would have to be processed on-board into the features required for the ML model. In developing hardware solutions for on-board satellite data processing, the implementation of higher order statistical tests may pose challenges. If these tests prove difficult to implement, exploring less optimal feature sets, or reducing compute and memory requirements through a data sampling scheme, could be a viable alternative. These features, initially chosen using recursive feature elimination from a broad range of options, highlight the need for close collaboration between the hardware implementation and machine learning teams. This collaboration is essential to find a balance between feature selection and the constraints of on-board hardware.

To enhance the throughput of BAQ based algorithms, several well-documented methods can be employed, albeit with some impact on the accuracy of the optimal quantizer. For instance, calculating the mean absolute value instead of standard deviation can speed up computations, as it avoids the need for slow square root calculations. Also, quantizing the allowable values of mean absolute and using a look-up table (LUT) for threshold and reconstruction values can significantly boost throughput. This approach also reduces the data size for transmission: using an 8-bit value for the LUT entry instead of a 32-bit floating point value for each mean absolute value. This not only improves throughput but also reduces the overhead of side information for each block, thereby increasing the compression ratio. These techniques are relevant for both BAQ and FFT-BAQ algorithms, given that FFT-BAQ is an extension of

BAQ with additional preprocessing steps. This strategy of modifying algorithmic approaches to meet hardware limitations is crucial for optimizing data processing in space missions.

Onboard inference could be implemented on FPGA using tools like Vitis AI. Now that we have shown that inference of SDNR in the image domain is possible from features derived from statistics in the raw data domain, a neural network (as a universal function approximator) could be trained to map these features to SDNR. Neural networks have a well-trodden path to implementation on hardware and there are toolchains available for further optimization of the onboard implementation e.g. weight quantization, connection pruning.

The output format could use an existing part of the CCSDS header to encode which compression algorithm has been used to encode each block. Something to highlight is that an algorithm chooser breaks the determinism that was there for downlink. The ground station no longer knows the exact size of the data to be downlinked before operations commence as there is no link between uplinked commands to make specific acquisitions and a determined compression rate (e.g. all acquired data compressed with BAQ4).

The models and algorithms developed under this Activity shall be further developed with collaboration with industry partners, with the aim being to demonstrate improvements on an in-orbit demonstrator mission.

The software implementation or future hardware implementation could be licensed to instrument developers. Collaboration with instrument developers would allow tuning for pre-processing and throughput requirements, new instruments to meet throughput constraints of onboard data buses, and tight integration with their interfaces and data formats. This would be in a runtime mode configuration.

A more realistic short-term goal is to integrate a software implementation or future hardware implementation with a platform manufacturer on a payload data processor, where the aim would be to reduce downlink reducing downlink constraints on the platform.

FUTURE WORK

The outcomes of this Activity lend themselves to a range of future work both as extensions of the work completed during the project and as new areas of investigation.

Offline Compression Implementation

An aim of the hardware roadmap development was to compare two approaches to the hardware architecture, runtime compression versus offline compression. NovaSAR-S was considered in the study as a potential future target hardware system and provided a good candidate for comparing a runtime compression mode with an offline compression operating mode, as it currently operates runtime compression.

Offline compression modes were found to be the preferred option, at least initially as they required the least modification of radar backend architectures like Airbus DS NIA Radar Backend (as used on NovaSAR). This future work would explore the interfaces to the High-Speed Data Recorder (HSDR) required in order to read and write from the inference module, and how this would impact both the mission concept of operations and the satellite and ground operations. Surrey Satellite Technology Ltd are already actively exploring offline on-board payload data processing architectures. One such project is the Incubed funded Flexible Intelligent Payload Chain (FIPC) where SSTL partnered with Craft Prospect Ltd and University of Surrey as application developers for an offline on-board processor designed to integrate the platform to process data from optical EO payloads. There is potential for this architecture to be extended to future SAR platforms.

Inference on Compressed Representation

The exploration of an offline processing mode architecture could be combined with exploring the idea of working with a BAQ compressed representation of the data directly (e.g. BAQ4 compressed) during inference. This approach would have advantages for systems where there is a coprocessor architecture with data already compressed by the radar backend. The task of the ML system would be to decide if data could be further compressed without loss of quality. For parts of the data where this was the case, the data would be decompressed and then recompressed with a lower bitrate algorithm. Areas of exploration include how the machine learning model would perform whilst working with features derived from the compressed data as input, and how this impacts the HSDR, as recompressed data would be smaller in size than the original.

Inference on Raw SAR for New Use Cases

The Activity developed a machine learning ready raw SAR dataset and a pipeline for generating further datasets. This could be adapted to other applications to enable rapid prototyping of new machine learning methods in the raw SAR domain. An initially targeted application could be a maritime application like ship detection and classification, where targets are sparse (ships) amongst the sea of data (the ocean). Much further development is expected in this area. In future work, the demonstrated algorithms could be embedded in a more realistic mission concept and demonstrated in simulation, with requirements preferably solicited from stakeholders with real use cases.

AI Compressor

A future project could explore the feasibility of a completely different approach to the problem of raw SAR data compression. This approach could be a pure learning approach, where the compression/decompression itself is learned by a model, rather than enhancing the use of existing compression algorithms by building on their design with a machine learning aspect.

Autoencoders are deep artificial neural networks, where the aim during training is to reproduce the encoder input at their output layer. A flavour of these architectures, 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 where some preliminary exploration has already been done for raw SAR data compression is using Vector Quantized Variational Autoencoders (VQ-VAE)[14], where discrete (rather than continuous) latent variables are used and distributions are categorical[15].

In this approach various VQ-VAE based model architectures could be explored, parameters and performance metrics tuned to reach these goals. The performance could be benchmarked against the state-of-the-art standard algorithms to evaluate the feasibility of AI as a solution to on-board SAR data compression. Within the model learning and evaluation process, the input would be raw SAR data and the ML model would aim for the highest quality approximation, with maximum compression ratio, as output. Outside of this process, the model output can be evaluated and verified against Single Look Complex (SLC) data to assess the quality of the product.

Neural Network Inference

The decision tree model selected in the course of this Activity is simple, effective and explainable. Using such a model is possible due to the features that can be derived from using knowledge of the statistics of the raw data. Now that we have proven that inference between the derived features and target variables can be successfully performed, we can look at other model types for on-boarding to satellite hardware. Neural networks can approximate any function so are expressive enough to learn the relationships here. They could also be used to learn (potentially more effective) features, rather than hand engineering features.

The tools available for implementing neural networks on hardware are fast advancing and provide multiple powerful optimizations like network weight quantization and architecture pruning to optimize hardware throughput against model inference performance.

CONCLUSION

The work carried out under SARAI demonstrated that the execution of machine learning models working on raw SAR data to perform predictions of reconstruction error in the image domain with various raw SAR data compression algorithms is feasible, albeit with a limited number of end-to-end solutions available when compatibility of all the elements is considered and many iterations of refinement required before an optimal solution can be found.

The achievements in the activity have applications across several ESA programmes. On-board processing to improve onboard data compression has many benefits to the EOEP (Earth Observation Envelope Program) in terms of eliminating or reducing data bottlenecks in EO missions, improving on-board data management, and reducing operational and infrastructure costs. This also benefits Science and Robotic Exploration missions (such as future interplanetary exploration missions with radar payloads), as better data compression can increase the possible scientific return of the mission. The activity also developed a machine learning ready raw SAR dataset and a pipeline for generating further datasets, which could be adapted to other applications to enable rapid prototyping of new machine learning methods in the raw SAR domain. Much further development is expected in this area. In future work, the demonstrated algorithms could be embedded in a more realistic mission concept and demonstrated in simulation, with requirements preferably solicited from stakeholders with real use cases. The suggestions of the Hardware Roadmap, created under this activity, should be further assessed, and followed to demonstrate the scheme on representative flight hardware. The demonstrator should explicitly address compliance with requirements on end benefits as defined by a relevant stakeholder. Completion of this work to address an end user application on representative flight hardware is estimated at €250k over fifteen months. This estimate is based on similar projects developing hardware testbeds with representative simulated scenarios representative of flight. A flight demonstration of the developed technology is a highly feasible mid-term goal. This is estimated at €2M total to develop flight models of an application demonstrator (ML technology fully integrated with a modified SAR instrument radar backend). This can then be integrated with a SmallSat bus and tested in-orbit. Clear requirements can be solicited from end users engaged through this and other parallel activities, resulting in a demonstration mission with measurable end benefits to these users. The compression demonstration would likely be part of a larger demonstrator mission. The work is estimated at eighteen months, using the TRL 6 milestone as a starting point.

References

1. R. L. Jordan, "The Seasat-A synthetic aperture radar system,", IEEE Journal of Oceanic Engineering, vol. 5, no. 2, pp. 154-164, April 1980.

2. V. Ignatenko, M. Nottingham, A. Radius, L. Lamentowski and D. Muff, "ICEYE Microsatellite SAR Constellation Status Update: Long Dwell Spotlight and Wide Swath Imaging Modes," 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, Brussels, Belgium, 2021.

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

4. E. Attema *et al.*, "Flexible Dynamic Block Adaptive Quantization for Sentinel-1 SAR Missions,", IEEE Geoscience and Remote Sensing Letters, vol. 7, no. 4, pp. 766-770, Oct. 2010.

5. P. Snoeij *et al.*, "Sentinel 1 – the future GMES C-band SAR mission", 2009.

6. C. Pieterse *et al.* "Metrics to evaluate compression algorithms for raw SAR data", IET Radar Sonar Navig, Mar 2019.

7. 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, vol. 36, no. 2, pp. 589-602, March 1998.

8. H. Saito, K. Ishimura, J. Hirokawa, T. Tomura, K. Ijichi, and K. Tanaka, 'Technical Challenges for Small SAR Satellites with High Performance', Proceedings of the AIAA/USU Conference on Small Satellites, 2021.

9. M. Cohen, A. Larkins, P. L. Semedo and G. Burbidge, "NovaSAR-S low cost spaceborne SAR payload design, development and deployment of a new benchmark in spaceborne radar," 2017 IEEE Radar Conference (RadarConf), Seattle, WA, USA, 2017

10. S. Doody, M. Cohen and J. Marquez-Martinez, "NIASAR-X Low Cost SAR Development," EUSAR 2018; 12th European Conference on Synthetic Aperture Radar, Aachen, Germany, 2018.

11. R. C. Huxtable, 'SAR RAW DATA COMPRESSION USING FPGA TECHNOLOGY', Master of Applied Science, Simon Fraser University, British Columbia, 2008.

12. CSIRO, "NovaSAR-1 User Guide", Available at https://research.csiro.au/cceo/novasar/novasar-introduction/novasar-1-user-guide/#modes

13. S. Sundaresan and F. Rivoallon, 'Analysis of Power Savings from Intelligent Clock Gating', 2012.

14. A. van den Oord, O. Vinyals, and K. Kavukcuoglu, 'Neural Discrete Representation Learning', Proceedings of the 31st International Conferene on Neural Information Processing Systems, Apr. 2017.

15. G. Pilikos, M. Azcueta, R. Camarero, and N. Floury, 'Raw Data Compression For Synthetic Aperture Radar Using Deep Learning', 8th International Workshop on On-board Payload Data Compression, Athens, Greece: ESA, 2022.