

SolarZip: An Efficient and Adaptive Compression Framework for solar EUV imaging data

Application to Solar Orbiter/EUI data

Zedong Liu^{1,2*}, Song Tan^{3,4*}, Alexander Warmuth³, Frédéric Schuller³, Yun Hong⁵, Wenjing Huang^{1,6}, Yida Gu^{1,6}, Bojing Zhu^{7,8,6}, Guangming Tan^{1,6}, and Dingwen Tao^{1,6**}

¹ Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China

² University of Electronic Science and Technology of China, Chengdu 610054, China

³ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany

⁴ Institut für Physik und Astronomie, Universität Potsdam, Karl-Liebknecht-Straße 24/25, 14476 Potsdam, Germany

⁵ Minzu University of China, Beijing 100081, China

⁶ University of Chinese Academy of Sciences, Beijing 100049, China

⁷ Yunnan Observatories, Chinese Academy of Sciences, Kunming 650216, China

⁸ Centre for Astronomical Mega-Science, Chinese Academy of Sciences, Beijing 100012, China

Received/ accepted

ABSTRACT

Context. With the advancement of solar physics research, next-generation solar space missions and ground-based telescopes face significant challenges in efficiently transmitting and/or storing large-scale observational data.

Aims. We develop an efficient compression and evaluation framework for solar EUV data, specifically optimized for Solar Orbiter Extreme Ultraviolet Imager (EUI) data, significantly reducing data volume while preserving scientific usability.

Methods. We systematically evaluated four error-bounded lossy compressors across two EUI datasets. However, the existing methods cannot perfectly handle the EUI datasets (with continuously changing distance and significant resolution differences). Motivated by this, we develop an adaptive hybrid compression strategy with optimized interpolation predictors. Moreover, we designed a two-stage evaluation framework integrating distortion analysis with downstream scientific workflows, ensuring that observational analysis is not affected at high compression ratios.

Results. Our framework SolarZip achieved up to 800× reduction for Full Sun Imager (FSI) data and 500× for High Resolution Imager (HRI_{EUV}) data. It significantly outperformed both traditional and advanced algorithms, achieving 3-50× higher compression ratios than traditional algorithms, surpassing the second-best algorithm by up to 30%. Simulation experiments verified that SolarZip can reduce data transmission time by up to 270× while ensuring the preservation of scientific usability.

Conclusions. The SolarZip framework significantly enhances solar observational data compression efficiency while preserving scientific usability by dynamically selecting optimal compression methods based on observational scenarios and user requirements. This provides a promising data management solution for deep space missions like Solar Orbiter.

Key words. Techniques: image processing – Methods: data analysis – Sun: corona – Space vehicles: instruments

1. Introduction

With the advancement of solar physics research, next-generation space solar missions and ground-based telescopes demand increasingly higher spatial and temporal resolution for observational data, making the efficient transmission and processing of large-scale data an urgent challenge. Solar Orbiter (Müller et al. 2020), a collaborative mission between the European Space Agency (ESA) and National Aeronautics and Space Administration (NASA), was successfully launched in February 2020. Its unique orbital design enables unprecedented close-up observations of the Sun (approaching as near as 0.28 AU) and provides the first high-resolution images of the solar polar regions. However, due to inherent telemetry constraints of deep space missions, the data observed by Solar Orbiter surpasses its transmission capabilities by a substantial margin, making efficient on-

board data compression essential for achieving the mission’s scientific objectives (Fischer et al. 2017). The Extreme Ultraviolet Imager (EUI, Rochus et al. (2020)), one of Solar Orbiter’s core remote sensing instruments, consists of a Full Sun Imager (FSI) and two High Resolution Imagers (HRI_{EUV} and HRI_{Ly α}), providing comprehensive observations from the chromosphere to the corona in EUV wavelengths (17.4 nm, 30.4 nm) and the Lyman- α band. EUI data exhibits significant high dynamic range characteristics, with intensity differences between bright and dark regions spanning several orders of magnitude. Furthermore, as Solar Orbiter’s orbital position and viewing angle continuously change, the characteristics of EUI data vary significantly, imposing strict adaptability requirements on compression algorithms.

Compression techniques for scientific data can be broadly categorized into lossless and lossy approaches (Patel et al. 2015; Di et al. 2024). While lossless compression guarantees exact data reconstruction, it typically achieves limited compression ratios of only 2:1 to 3:1 for scientific floating-point data (Patel

* These authors contributed equally to this work.

** Corresponding author; taodingwen@ict.ac.cn

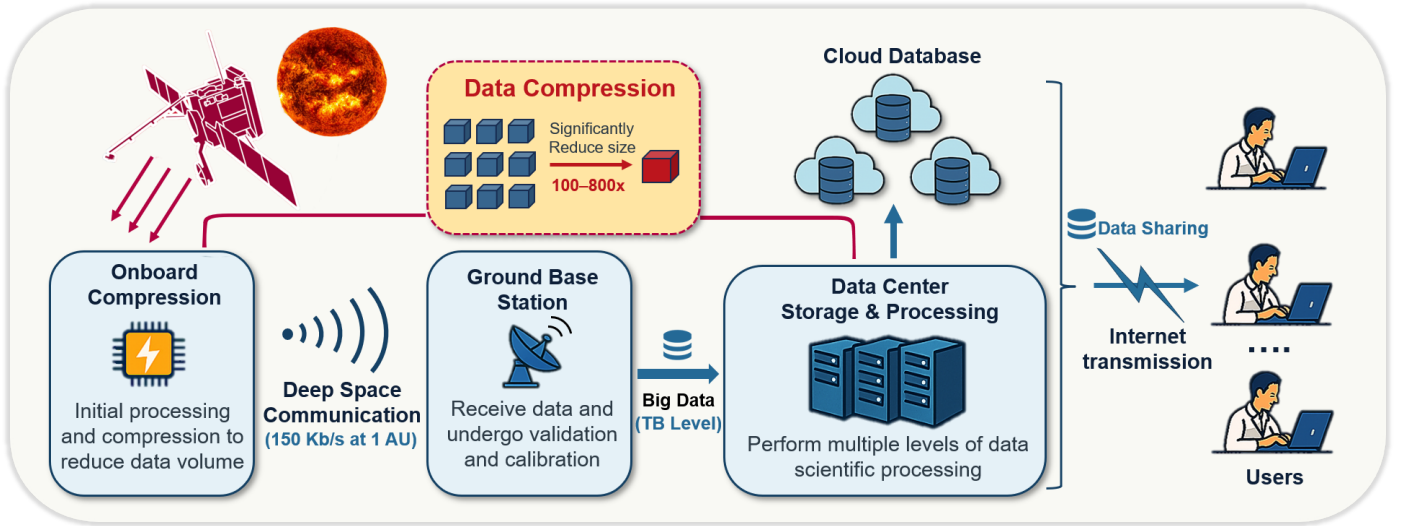


Fig. 1: Lifecycle of Solar Orbiter observation data. This figure illustrates the application of data compression techniques throughout the lifecycle of solar observation data. These techniques significantly reduce data volume, addressing communication and storage challenges, particularly in onboard systems and data centers.

et al. 2015). In contrast, error-bounded lossy compression can achieve much higher ratios while maintaining scientific usability by controlling data distortion within user-specified tolerances (Cappello et al. 2019; Di et al. 2024).

In recent years, a new generation of lossy compressors designed specifically for scientific data has emerged, including SZ (Di & Cappello 2016; Tao et al. 2017; Liang et al. 2018b; Zhao et al. 2021), ZFP (Lindstrom 2014), MGARD (Ainsworth et al. 2019), and SPERR (Li et al. 2023). Unlike traditional lossy compressors such as JPEG (Wallace 1991; Taubman & Marcellin 2002), these error-bounded lossy compressors are designed to compress scientific data while providing strict error control based on user requirements. These compressors have been successfully applied across various scientific domains. In climate simulation, Baker et al. (2014, 2016, 2017, 2019) employed lossy compression on data produced by the Community Earth System Model. For cosmological simulations, (Jin et al. 2020; Pulido et al. 2019; Jin et al. 2021) proposed efficient schemes to reduce storage and transmission costs for Nyx and WarpX simulation codes. In astronomical observations, studies have evaluated how transform-based algorithms affect radio astronomy data quality (Peters & Kitaeff 2014; Vohl et al. 2015; Chege et al. 2024), other researchers explored efficient onboard compression algorithms for satellite missions using Cassini observational data (Xie et al. 2021; Zhang et al. 2025).

Traditional compression methods have been explored for solar data. The RICE encoding algorithm (Rice & Plaunt 1971), which relies on basic preprocessing and encoding, achieves a maximum compression ratio of 20× for Solar Orbiter EUV data (result from Section 3). Fischer et al. (2017) implemented JPEG2000, a wavelet-transform-based compression method, yet its compression ratio was limited to 30×, with significant quality degradation at higher compression levels. Recent efforts have explored machine learning approaches (Zafari et al. 2022, 2023; Liu et al. 2024a,b), such as Attention Mechanisms and Generative Adversarial Networks (GANs), achieving promising compression ratios but introducing substantial training and computational overhead.

Despite these advances, existing approaches for solar data compression have notable limitations: (1) They rely primarily

on traditional lossy compression algorithms with relatively naive strategies, achieving limited compression ratios; (2) No comprehensive study has systematically applied or evaluated advanced error-bounded lossy compression techniques on solar EUV data; and (3) Previous work lacks interdisciplinary insights from both astronomy and computer science, failing to comprehensively illustrate the impact of lossy compression on scientific analyses of solar observations.

To address these gaps, this paper introduces SolarZip, an efficient compression and evaluation framework for solar EUV imaging data, with the following contributions:

- We analyze four advanced lossy compressors across two Solar Orbiter EUV datasets with 14 settings. They demonstrate clear advantages over traditional methods, but still face limitations.
- We introduce SolarZip, a comprehensive compression and evaluation framework for solar EUV imaging data, particularly optimized for Solar Orbiter/EUV data.
- We design a two-stage data evaluation framework that integrates strict error control with downstream scientific workflows: data distortion analysis and scientific post-hoc analysis. This ensures that compressed data remains suitable for critical scientific research.
- We develop an adaptive hybrid compression strategy with optimized predictors to enhance compression quality. This method dynamically selects optimal compression methods based on different observational scenarios and user requirements, achieving up to 800× reduction for FSI data and 500× for HRI_{EUV} data.

This paper is structured as follows: Section 2 introduces the fundamentals of data compression techniques. Section 3 presents a comprehensive experimental comparison between advanced compression algorithms and traditional methods. In Section 4, we detail the SolarZip framework, including its compression workflow, hybrid strategy, and optimization methodology. Section 5 provides a thorough evaluation of the SolarZip framework, demonstrating its superior performance through extensive experimental results. Finally, Section 6 concludes the paper and outlines promising directions for future research.

2. Data compression foundation

There are some traditional lossy compressors for images and videos such as JPEG and MPEG, but they do not have a good performance on the scientific data. Error-bounded lossy compression is a new generation that are designed for scientific data.

2.1. Error-bounded lossy compression

Generally, error-bounded lossy compressors require users to set an error type, such as the point-wise absolute error bound and point-wise relative error bound, and an error bound level (e.g. $\text{abs error} = 10^{-3}$). The compressor ensures that the differences between the original data and the reconstructed data remain within the user-set error tolerance (Cappello et al. 2019). This ensures that the reconstructed data maintains a controlled level of accuracy, making it suitable for scenarios where precision is paramount.

The workflow of error-bounded lossy compressors can be summarized in the following steps: (1) Data preprocessing, such as domain transformation and data blocking. (2) Decorrelation via compression models, which are broadly categorized as either prediction-based or transform-based. (3) Quantization with controlled error to achieve data compression. (4) Further lossless compression of quantized codes and other parameters using techniques like arithmetic coding. The core component of a lossy compressor is its decorrelation model, as it critically influences both compression efficiency and speed.

The SZ family of compressors are representative examples of prediction-based compression models. Their predictors include linear regression predictors, Lorenzo predictor (used in SZ1.4–SZ2.0, (Tao et al. 2017; Liang et al. 2018b)), and the spline interpolation predictor (introduced in SZ3, (Zhao et al. 2021; Liang et al. 2022)). Previous studies (Tao et al. 2019; Liang et al. 2019; Zhao et al. 2020) have explored the strengths and limitations of these predictors, as well as their cooperative use. Overall, the SZ series achieves high compression ratios while maintaining a favorable compression speed, making it a versatile solution for scientific data compression.

Conversely, transform-based compression models employ different techniques to achieve decorrelation. For instance, ZFP (Lindstrom 2014) utilizes near-orthogonal transforms, while SPERR (Li et al. 2023) applies the CDF 9/7 biorthogonal wavelet transform. ZFP is particularly notable for its high-speed performance; however, its limitation is that the compression ratio is constrained. The advantage of SPERR is that the hierarchical multidimension DWT in SPERR can effectively capture the relevance between data points, which brings a high compression ratio after the SPECK encoding. One limitation of SPERR is that the transform and encoding processes have high computational costs, and hence its execution speed is relatively low, typically around 30% of SZ3.

2.2. Applications of compression in solar observation data.

Using Solar Orbiter as an example (fig. 1), we explain how data compression techniques are integrated into the lifecycle of solar observational data. The process begins with data collection by the spacecraft's remote-sensing and in-situ instruments. Once generated, the data undergoes initial processing and compression onboard. Since space communication bandwidth is limited (nominal 150 kbit/s at 1 AU) and storage is constrained, compression is essential to facilitate efficient transmission. The data is downlinked via X-band telemetry to ESA's deep-space ground

stations, where it is received and undergoes initial processing. Upon reaching ground stations, it undergoes initial validation and calibration. Subsequently, the data is transferred to dedicated scientific data centers, where scientists and researchers conduct various levels of data processing. Compression may also be applied at data centers to manage storage and transmission challenges posed by the vast volume of data. Finally, publicly available datasets are released through mission archives, enabling broader scientific access.

2.3. Problem and metrics description

In this paper, we focus on the design and implementation of a lossy compression algorithm for solar EUV observational data represented by Solar Orbiter/EUI. The key goal is to achieve efficient data compression, significantly reducing data volume. At the same time, it is crucial to ensure that the quality of the reconstructed data meets the requirements of scientific research. To achieve a better compression methods, we employ the following metrics, which are widely used in prior literature and considered standard in the field (Leung & Apperley 1994).

Metric 1 CR: We use the compression ratio (CR) to measure the compression performance. The Compression Ratio (CR) is calculated by dividing the original data size by the compressed data size. A CR of 100 indicates that the compressed data is 1/100 the size of the original.

$$CR = \frac{\text{Original Size}}{\text{Compressed Size}} \quad (1)$$

Metric 2 PSNR: For data point i , let $e_{\text{abs}_i} = x_i - \tilde{x}_i$, where $[e_{\text{abs}}]$ is the absolute error. Also, we denote the range of X by R_X . To evaluate the average error in the compression, we first use the popular root mean squared error (RMSE).

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (e_{\text{abs}_i})^2} \quad (2)$$

The peak signal-to-noise ratio (PSNR) is another commonly used average error metric for evaluating a lossy compression method (Berger 2003), especially in visualization. A higher value of PSNR represents less error. It is calculated as follows.

$$PSNR = 20 \cdot \log_{10} \left(\frac{R_X}{RMSE} \right) \quad (3)$$

Metric 3 ρ : The Pearson correlation coefficient ρ is employed to assess the linear correlation between the original and reconstructed datasets. A correlation coefficient of at least 0.9999 is generally required to ensure high fidelity.

$$\rho = \frac{\text{cov}(X, \tilde{X})}{\sigma_X \sigma_{\tilde{X}}} \quad (4)$$

Metrics 4 SSIM: The Structural Similarity Index (SSIM) is another popular metric for evaluating the perceptual quality of images in compression. Unlike traditional error-based metrics, SSIM considers structural information, including luminance and contrast, which better reflects human visual perception. A higher SSIM value corresponds to greater similarity, with a value of 1 signifying perfect structural and perceptual equivalence.

$$SSIM(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (5)$$

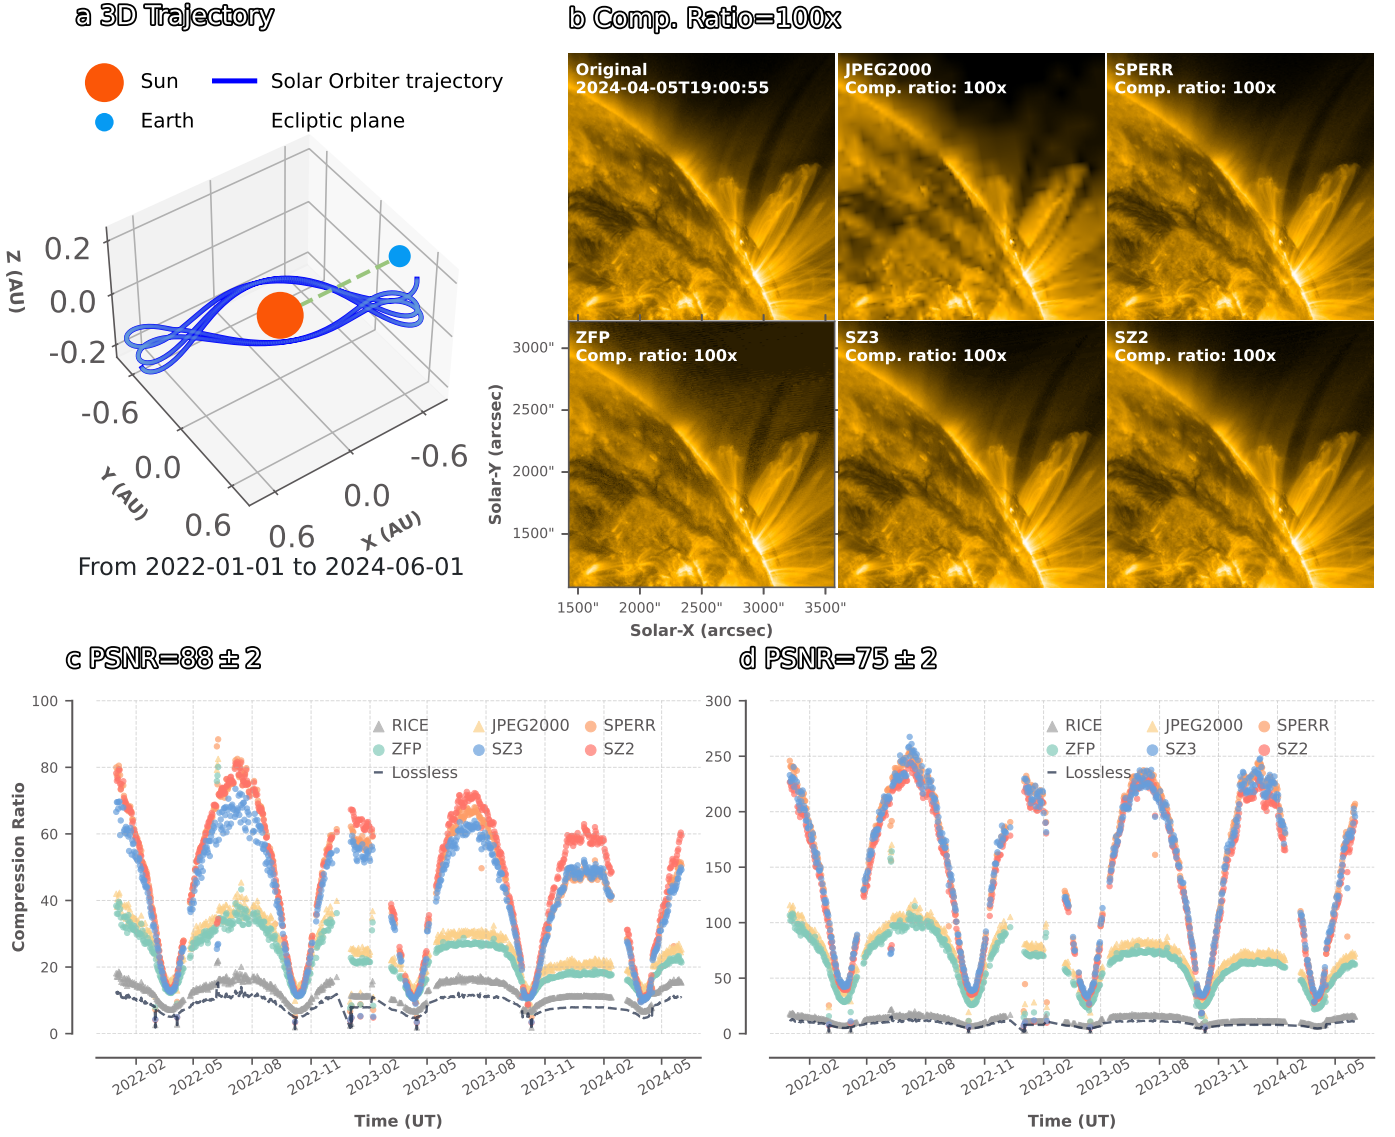


Fig. 2: Panel a: Visualization of the Solar Orbiter’s trajectory based on the 2.5-year FSI dataset. Panel b: Comparison between the original image acquired on April 5, 2024, and the reconstructed images from five compression algorithms, at the same compression ratio of 100x. Panel c: Compression ratio trends over time at a fixed quality level corresponding to PSNR = 88. Panel d: Compression ratio trends over time at a fixed quality level corresponding to PSNR = 75. Color coding indicates different compression methods. RICE and lossless compression (GZip) have fixed compression ratios and are used as baselines for reference. Advanced lossy compressors demonstrate significant advantages in compression performance.

Post-Hoc Metrics: In addition to conventional data distortion metrics, this paper also introduces domain-specific scientific downstream analysis metrics. Aiming for a comprehensive evaluation of the impact of compression on EUV data, we provide a detailed discussion in Section 4.4.

3. Comparison of advanced and traditional compression algorithms

In this section, we evaluate the performance of advanced lossy compression methods and traditional image compression techniques on EUV data. Based on the metrics specified in Section 2.3, we perform a comparative analysis to assess the strengths and weaknesses of both classes of compressors.

3.1. Experimental setup

Environment: The advanced compression methods (error-bounded lossy compressor) included in the experiment are SZ2, SZ3, ZFP, and SPERR, while the traditional image compression techniques include JPEG2000 and RICE (native algorithm of the FITS format). As a reference, we also included the lossless compression algorithm GZip. The experimental environment is deployed on a cloud server equipped with two 16-core Intel Xeon Gold 6151 3.00GHz CPU, and 1,007 GB of RAM.

Our test EUV dataset consists of two parts:

FSI Dataset: EUV/FSI is an EUV imager observing the full solar disk in the 17.4 nm and 30.4 nm EUV passbands. It continually provides synoptic observations with a variable cadence depending on telemetry allocations and observing mode. We adopted a daily sampling strategy, selecting 17.4 nm images

from the latest EUI data release 6.0 (Kraaikamp et al. 2023). During the data screening process, we excluded observations when FSI was operating in coronagraph mode to ensure data consistency and representativeness. The final FSI dataset totals 5.8 GB, covering observations from different orbital positions throughout Solar Orbiter’s nearly 2.5 years from January 2022 to June 2024, comprehensively reflecting the characteristic variations in FSI data.

HRI_{EUV} Dataset: The EUI/HRI_{EUV} dataset is based on a flare observation campaign conducted on April 5, 2024. During this campaign, HRI_{EUV} continuously observed a solar limb active region for approximately 4 hours from 19:59 to 23:59 with a temporal cadence of 16 seconds. This long time series observation generated a substantial amount of high-resolution data, totaling 4.2 GB. This dataset is particularly suitable for evaluating the performance of compression algorithms when processing complex, rapidly evolving solar active regions.

All data used in this work are level-2 products, which have undergone preliminary standard calibration processing. Notably, the publicly released EUI datasets are natively compressed using the RICE algorithm. Although this native compression does not affect level-1 data, it introduces minor artifacts in level-2 data, potentially reducing compression performance. Despite this, we choose level-2 data due to its calibration, which enhances reliability for post-hoc scientific analysis. The combination of these two datasets provided us with a comprehensive opportunity to evaluate the performance of the SolarZip framework under different observational modes, including full-disk low-resolution observations and local high-resolution dynamic feature observations.

3.2. Performance analysis and comparison

To ensure a rigorous and objective comparison of compression performance across different algorithms, a unified quality metric PSNR (as equation 3) is adopted. A higher compression ratio at a given image quality indicates a superior algorithm. Fig. 2 c shows the compression ratios of different algorithms over observation time at PSNR = 88. The SZ family and SPERR demonstrate the highest compression ratios, while JPEG2000 and ZFP achieve similar but inferior results compared to the best-performing algorithms. The RICE algorithm (native algorithm of the FITS format) exhibits the lowest compression ratio. Fig. 2 d shows the results at PSNR = 75, where the performance gap widens further. The best-performing algorithms achieve compression ratios close to 300×, whereas JPEG2000 is limited to 50×–100×. As an example, on June 1, 2023, SZ3 achieved a compression ratio 2.4× higher than JPEG2000 and 18.4× higher than RICE.

Figure 2 b compares the reconstructed images of four advanced lossy compressors and JPEG2000 on the sample data. At the same compression ratio, JPEG2000 exhibits noticeable distortion and severe compression artifacts in many regions. In contrast, the four advanced compressors deliver superior performance. While ZFP shows minor deviations, the reconstructed images from SZ3, SZ2, and SPERR are visually indistinguishable from the original, maintaining excellent image quality.

Strength: Experiments on full time series dataset and corresponding visual analyses confirm the superior performance of advanced compression algorithms on EUI data. Quantitative evaluation reveals that significantly outperform traditional methods: achieving 10–30× higher compression ratios over RICE and 1.5–3× over JPEG2000, under equal reconstruction quality. This breakthrough is primarily attributed to the error-bound control

mechanisms employed by these algorithms, which enable efficient data reduction while preserving image fidelity.

3.3. Observation and motivation

However, despite advanced lossy compressors show clear advantages over traditional algorithms, we also observed their inherent limitations, which motivated the design of SolarZip.

Observation 1: Compression ratios vary over time and closely follow changes in the spacecraft’s distance from the Sun. Specifically, when the spacecraft is farther from the Sun, the compression ratio increases under the same error bound, when closer, the compression ratio decreases. This can be explained by the reduced apparent size of the Sun in distant images, leading to a more uniform background distribution that is easier to compress.

Observation 2: No single compressor performs best for EUI data. As shown in Fig. 2, at PSNR of 75 dB, ZFP exhibits the lowest compression efficiency, while SZ3 and SPERR performing similarly better. At a higher PSNR threshold of 88 dB, the compression performance analysis reveals that SZ2 achieves superior results compared to other compressors, while SZ3 exhibits marginally lower performance than SPERR. We further investigate the reasons behind these variations. SZ3 relies on global interpolation and Lorenzo predictors, which perform well on datasets with strong global continuity. In contrast, SZ2 utilizes a block-based Lorenzo predictor, achieving higher accuracy for locally continuous data.

Motivation: These observations confirm our key insight: EUI data characteristics are highly dependent on observation conditions, making it impossible to define a single optimal compression strategy. Therefore, a dynamic multi-compression framework is necessary to adapt to complex observational images and meet diverse scientific requirements.

4. SolarZip compression framework

As discussed above, designing a comprehensive compression framework for EUI data must address the diverse characteristics variability inherent in the data. This requires the dynamic selection of compression strategies tailored to varying data characteristics, while maintaining high compression efficiency. Thus, we propose the SolarZip data compression and evaluation framework (Fig. 3).

4.1. Overview of SolarZip

The SolarZip workflow consists of four stages: (1) Initialization: pre-processing FITS files and setting configurations. (2) Compression Stage: Selecting the optimal compression strategy, automatically optimizing it, and then applying the compressor to the input data. (3) Distortion Analysis: Analyzing errors, evaluating distortion, and visualizing results. (4) Post-hoc Analysis: Performing downstream tasks such as Coronal Structure analysis and dynamic feature analysis on reconstructed data.

Our workflow follows a modular design, enabling the efficient parallel compression of large volumes of FITS files. Users need to provide a configuration file and the FITS files. During the initialization stage, the system is configured based on the user’s configuration parameters, and the FITS files are preprocessed. The preprocessing includes splitting the HDUs in the FITS files and extracting the data to be compressed. The data is then converted into binary format and passed to the next phase. FITS files

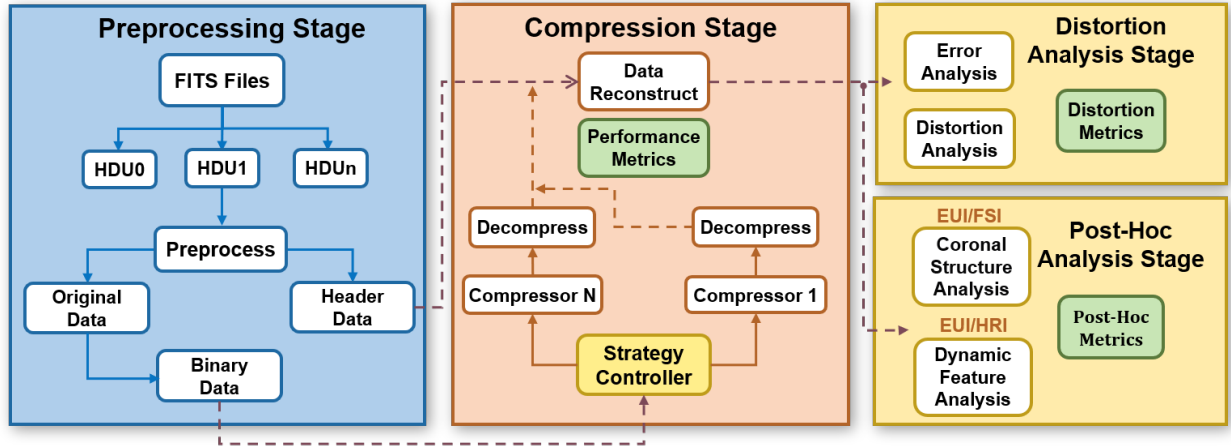


Fig. 3: Overview of SolarZip framework. The system consists of three stages: preprocessing, compression, and analysis. The core algorithmic innovation lies in the Strategy Controller can automatically tune and select the optimal compression strategy. The subsequent two stages of comprehensive analysis ensure that the decompressed data remain suitable for scientific purposes.

are processed in batches, with the compression of each file occurring in parallel. Importantly, due to the adaptive strategy controller, we achieve the highest compression ratio for images from different observational scenarios and user requirements. Details about the adaptive strategy and optimization are presented in the next subsection.

After decompression, distortion analysis is performed by comparing the decompressed data with the original, producing distortion metrics. The data processor then reconstructs the decompressed data back into FITS format, and the reconstructed FITS files are used for further post-hoc analysis.

4.2. Adaptive hybrid compression strategy

We conducted a thorough evaluation of four lossy compressors on Solar Orbiter/EUI data (detailed in Section 3). The results indicate that no single compressor consistently outperforms the others across various observational scenarios and scientific objectives. The core reason behind these variations lies in the distinct data prediction and transformation mechanisms of these advanced lossy compression algorithms (elaborated in Subsection 2.1).

In this subsection, we propose an adaptive optimization strategy that automatically selects the most suitable compression algorithm based on data characteristics in different scenarios. This strategy leverages a sampling-based approach to dynamically determine the optimal parameter configuration.

Our adaptive hybrid compression strategy is depicted in Fig. 4. We categorize the precision requirements of EUI data into two types: relaxed and strict, using a relative error threshold of 1×10^{-4} as the boundary. For a relaxed error bound ($eb > 1 \times 10^{-4}$), the optimal strategy is chosen between spline interpolation prediction and Lorenzo prediction. Under a strict error bound ($eb \leq 1 \times 10^{-4}$), we select the best strategy from transform-based predictors and linear regression predictors.

Specifically, we employ a heuristic sampling approach, where N data points are uniformly selected from the dataset to guide the compression strategy selection. For a relaxed error bound ($eb > 1 \times 10^{-4}$), the tuning process consists of four steps: (1) Uniformly sampling N data points from the dataset. (2) Optimizing the spline interpolation predictor (Trial run S1 in Fig. 4) by selecting the best-fit interpolation method (linear or

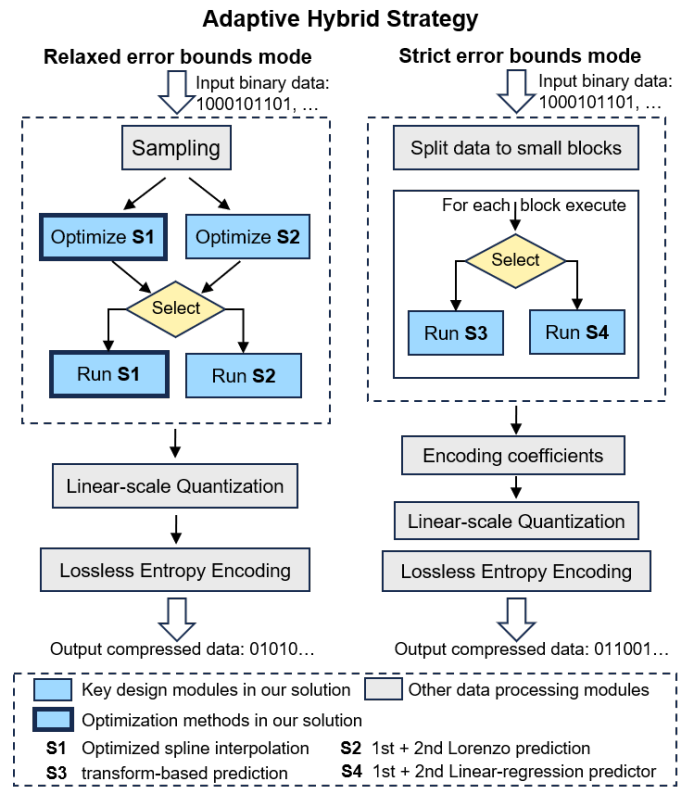


Fig. 4: Steps of Adaptive Hybrid Compression Strategy. The left panel shows the compression strategy under loose error bounds while the right panel under strict error bounds. Our system dynamically selects the optimal compression strategy and optimize it. The different strategies are denoted by S1–S4 in the figure.

cubic) and optimizing the sequence of interpolation dimensions. (3) Optimizing the Lorenzo predictor by dynamically selecting between first-order and second-order predictions. (4) Selecting the best strategy with the highest compression ratio. For strict error bounds ($eb \leq 1 \times 10^{-4}$), the tuning process follows the same principle, except that run A corresponds to the transform-based predictor, while run B corresponds to the regression-based predictor.

Linear regression prediction and Lorenzo prediction have been effectively applied in previous studies, both of which are block-based prediction methods. In this work, we conduct an offline analysis to determine the optimal block size of 8×8 , which is then set as the default parameter in our framework. Both linear regression and Lorenzo predictors support first-order and second-order prediction functions. Following the approach proposed in previous studies (Zhao et al. 2020), we dynamically select between first-order and second-order prediction functions for each data block. This ensures that every data block applies the optimal compression strategy. We apply the optimal Orthogonal Transformation in ZFP (Lindstrom 2014) as our prediction model, because its de-correlation efficiency has been shown to be more effective than that of other transforms, such as the discrete cosine transform or wavelet transform.

4.3. Optimization for spline interpolation predictor

The compression method based on classical spline interpolation can achieve a high compression ratio under large error bounds. However, in some cases, noticeable image quality degradation occurs. For instance, in highly non-smooth regions of solar EUV images, such as flares, the interpolation predictor introduces visible compression artifacts in those areas. This issue arises because the basic interpolation-based predictor suffers from considerably low accuracy in long-range interpolation (Liu et al. (2022)). Since it does not control the maximum stride length, the prediction accuracy becomes fairly low when the interpolation spans a long distance in the data array. To address these problems, SolarZip implements two key optimizations:

Grid-wise anchor points interpolation. In the interpolation process, we specifically introduce grid-wise anchor points. Anchor points are pre-known data points, which are losslessly encoded and stored during compression. These anchor points divide the entire dataset into multiple blocks, and all other data points are predicted using points within a certain range, employing a multi-level interpolation method. This method effectively addresses the issues caused by long-range predictions. It is noteworthy that we found that if an appropriate stride is set for the anchor grid, the overhead associated with storing the losslessly compressed anchor points becomes negligible.

Level-wise interpolation with error bound auto-tuning. We set different error bounds at different levels of interpolation (as opposed to the unified error bounds used in SZ3). In our two-dimensional data, 75% of the data points fall within the lowest interpolation level (Level 1), which are predicted by higher-level reconstructed data points, while the remaining 25% of the data points are predicted at higher levels. Therefore, setting smaller error bounds at higher levels helps ensure overall prediction accuracy, thus improving compression quality.

$$e_l = \frac{e}{\min(\alpha^{l-1}, \beta)} \quad (\alpha \geq 1 \text{ and } \beta \geq 1) \quad (6)$$

The level-wise error bounds e_l are dynamically adjusted based on equation 6. The parameters α and β are introduced, where e represents the global error bound set by the user. We perform offline testing with parameter sets $\alpha = \{1, 1.5, 2\}$ and $\beta = \{2, 3, 4\}$, comparing the bit-rate and PSNR values across different parameter configurations. Ultimately, we select $\alpha = 1.5$ and $\beta = 4$ as our optimal parameters.

4.4. Post-hoc analysis stage

In addition to standard distortion analysis, we implement specialized post-hoc analysis tailored to the scientific requirements of solar physics research.

4.4.1. FSI large-scale coronal structure analysis

For the FSI data, we focus on evaluating how compression affects large-scale dynamic structures in the solar corona. We implement a circular intensity extraction method, where a virtual circle is placed at 1.05 solar radii from the disk center, corresponding to the lower corona region where many important dynamic phenomena occur. Intensity values are sampled along this circle to generate a one-dimensional intensity profile that captures the coronal structures. The intensity profiles from the original and reconstructed compressed images are then compared to assess how different compression algorithms preserve the coronal features. We compare several metrics, including:

- Morphology-based visual comparison of full-disk FSI images
- Intensity profile correlation coefficient

This analysis is particularly important for studying large-scale coronal evolution and identifying the onset of coronal mass ejections, where subtle changes in intensity distribution can have significant scientific implications.

4.4.2. HRI_{EUV} small-scale dynamic feature analysis

For the HRI_{EUV} data, our post-hoc analysis focuses on the preservation of small-scale dynamic structures in selected frames. We examine features such as plasma flows and fine magnetic structures, which are crucial for understanding energy transport and release in the solar atmosphere. We implement feature tracking algorithms to identify and characterize dynamic features in both the original and compressed reconstructed data. The analysis includes:

- Morphology-based visual comparison of local dynamic features in HRI_{EUV} images
- Difference maps between original and compressed images
- Median and standard deviation of the pixel difference distribution in the difference maps

In contrast to FSI's analysis of large-scale structures, for HRI_{EUV}, we focus on a detailed analysis of representative small-scale structures, comparing the sensitivity of different features under various compression ratios to propose targeted dynamic compression strategies.

5. Evaluation results

We tested the compression performance and reconstruction quality of SolarZip alongside five advanced lossy compression techniques on EUV data. Our evaluation framework consists of two analytical stages: The first stage calculates compression performance metrics by comparing the original and reconstructed data, and the second stage incorporates expert domain knowledge for systematic post-hoc analysis, ensuring a comprehensive multi-dimensional assessment of the compression evaluation. Additionally, we conducted simulation experiments based on the hardware conditions of Solar Orbiter/EUI, verifying the significant improvement in data transmission efficiency enabled by lossy compression algorithms.

Table 1: Comparing four compressors on FSI and HRI_{EUV} datasets under three error bounds. Key metrics include compression ratio (higher=better), PSNR (higher=better), correlation coefficient (1.0=perfect), and SSIM (1.0=perfect).

| Error-bound | Compressor | FSI-Dataset | | | | HRI _{EUV} -Dataset | | | |
|-------------|-----------------|---------------|--------|-------------|------|-----------------------------|-------|-------------|-------|
| | | Comp Ratio | PSNR | Coefficient | SSIM | Comp Ratio | PSNR | Coefficient | SSIM |
| $1e-4$ | SPERR | 43.46 | 91.74 | 1.00000 | 0.74 | 5.69 | 90.78 | 1.00000 | 0.99 |
| | ZFP | 32.48 | 100.84 | 1.00000 | 0.84 | 5.83 | 103.2 | 1.00000 | 1.00 |
| | SZ3 | 40.24 | 89.29 | 1.00000 | 0.76 | 5.14 | 85.33 | 1.00000 | 0.99 |
| | SZ2 | 45.63 | 90.07 | 1.00000 | 0.77 | 4.83 | 84.79 | 1.00000 | 1.00 |
| | SolarZip | 56.28 | 92.88 | 1.00000 | 0.76 | 5.71 | 90.14 | 1.00000 | 0.99 |
| $1e-3$ | SPERR | 151.98 | 77.23 | 0.99999 | 0.45 | 10.51 | 65.57 | 0.99997 | 0.96 |
| | ZFP | 62.48 | 82.33 | 0.99999 | 0.58 | 10.75 | 73.08 | 0.99999 | 0.99 |
| | SZ3 | 149.55 | 73.99 | 0.99998 | 0.46 | 10.40 | 65.04 | 0.99997 | 0.96 |
| | SZ2 | 143.72 | 74.91 | 0.99998 | 0.56 | 10.28 | 64.79 | 0.99996 | 0.96 |
| | SolarZip | 153.68 | 76.54 | 0.99999 | 0.46 | 11.04 | 69.75 | 0.99998 | 0.97 |
| $1e-2$ | SPERR | 660.23 | 61.72 | 0.99981 | 0.39 | 112.53 | 49.98 | 0.99928 | 0.39 |
| | ZFP | 210.7 | 68.20 | 0.99987 | 0.46 | 16.17 | 61.12 | 0.99993 | 0.816 |
| | SZ3 | 684.11 | 57.17 | 0.99952 | 0.42 | 79.56 | 48.37 | 0.99869 | 0.38 |
| | SZ2 | 557.26 | 57.87 | 0.99947 | 0.39 | 82.26 | 46.08 | 0.99897 | 0.35 |
| | SolarZip | 695.31 | 60.92 | 0.99984 | 0.41 | 115.78 | 57.80 | 0.99974 | 0.34 |

This detailed examination evaluates whether the compressed data retains sufficient information to support the scientific analysis of transient solar phenomena and small-scale structures, which is essential for studies of magnetic reconnection, wave dynamics, and plasma heating mechanisms in the solar atmosphere. By combining these specialized post-hoc metrics with standard distortion analysis, we provide a comprehensive evaluation of compression performance that directly addresses the scientific use cases for EUV data.

5.1. Evaluation of compression performance

Table 1 shows the test results of SolarZip against four error-bounded lossy compression algorithms (ZFP, SZ3, SZ2, SPERR) on the FSI dataset and HRI_{EUV} dataset. The experiments show that SolarZip achieves optimal compression ratios on both datasets, with an improvement of 5.6-30.4% over the second-best compressor, while maintaining excellent fidelity (PSNR > 60 dB). This advantage is mainly due to our proposed Adaptive Hybrid Compression Strategy (AHCS), which demonstrates excellent adaptability to the data and is able to optimize and select the optimal compression strategy, as well as the optimization of the spline interpolation predictor, which improves the quality of the reconstructed images under high error bounds. Thus, SolarZip achieves a high compression ratio for solar science data while ensuring scientific usability.

To enhance the rigor and objectivity of the compression performance comparison, we plot the rate-distortion curve for our solution and other lossy compressors, comparing the distortion quality at the same rate. Here, rate refers to the bit rate in bits per value, and we use the peak signal-to-noise ratio (PSNR) to measure the distortion quality. PSNR is calculated by equation 3 in decibels. Generally speaking, in a rate-distortion curve (Berger 2003), a higher bit rate indicates that more bits are required to store each value, resulting in higher quality of the reconstructed data after decompression, as reflected by a higher PSNR.

As discussed in Section 4.2, we designed an adaptive hybrid strategy to optimize the compression quality across the entire bit-rate range. Fig. 5a presents the rate-distortion curves of our algorithm compared to five other methods on the FSI dataset. The

results demonstrate that our adaptive hybrid strategy plays a crucial role in improving compression quality. As shown in Fig. 5a, our compression algorithm achieves near-optimal quality across almost all bit rates. Particularly, for bit rates below 1.0, our method exhibits notably superior compression quality compared to SZ3 and SZ2, attributed to its dynamic selection between the automatically optimized spline interpolation predictor and the Lorenzo predictor. Moreover, our approach performs comparably to the SPERR compressor, which is specifically designed for high-quality compression but is unsuitable for our task due to its inherent limitations, as will be discussed later. At a PSNR of 77, our method achieves a compression ratio 300% higher than JPEG2000, demonstrating the efficacy of our spline interpolation predictor optimizations. When the bit rate exceeds 1.0, our solution surpasses SPERR, becoming the most rate-distortion efficient compressor. This is because, under a relaxed error bound greater than 1×10^{-4} , our strategy accurately selects the linear regression predictor and the transformation-based predictor (using the same orthogonal transformation matrix as ZFP). Notably, our method consistently outperforms JPEG2000, a widely used compression standard in astronomy.

Fig. 5b presents the rate-distortion curves of all compressors on the HRI_{EUV} dataset. The results demonstrate that our solution achieves the best performance among all six compressors. It can be observed that at a bit rate of approximately 3, our rate-distortion curve exhibits a distinct inflection point while maintaining superior quality. This behavior results from an adaptive adjustment in our selection strategy as the error bounds transition from a strict to a relaxed mode, with the threshold determined through offline analysis. Additionally, at the same PSNR level (e.g., PSNR = 60), our method achieves a 50% higher compression ratio than JPEG2000, further validating the effectiveness of our optimizations on the spline interpolation predictor in improving image quality.

We conducted a simulation experiment on the compression and data transmission process aboard the Solar Orbiter. Solar Orbiter/EUI is powered by a LEON3FT UT699 66 MHz single-core CPU, delivering an average telemetry data rate of around 300 Kbit/s (Marinrodriga et al. 2021). We simulated the data transmission process after compression and recorded the total




Fig. 5: Rate-distortion curves on different datasets. Different compressors are distinguished by color, with our method indicated by the red line. A higher PSNR corresponds to better image quality at the same bit rate. SolarZip demonstrates the best overall performance on both datasets.




Fig. 6: Comparison of elapsed transmission times for different compressor with a data volume of 1000 MB. The elapsed time includes compression time (red) and data transmission time (blue), reflecting the end-to-end transmission efficiency. Uncompressed transmission takes about 7.5 hours, while SolarZip demonstrates the highest efficiency, reducing the total time by 270 \times .

time, including both compression and transmission durations. Fig. 6 presents the results under two different error bounds. Our solution achieved the shortest elapsed time, demonstrating the highest data transmission efficiency. For instance, transferring 1000 MB of uncompressed data takes approximately 7.5 hours, whereas our method completes the task in just 101 seconds—improving efficiency by a factor of 270. This is due to the superior compression performance and high compression speed of our method. Although SPERR achieved a slightly higher compression ratio than our solution, its slow compression speed resulted in a significantly longer total time.

5.2. Evaluation of post-hoc analysis

After evaluating generic compression performance, our focus shifts to whether reconstructed FSI and HRI_{EUV} images can meet the requirements of solar physics observational research. Based on the data’s inherent variations and specific research content,

we propose what we consider appropriate dynamic compression strategies.

5.2.1. FSI reconstructed image analysis

As shown in Fig. 7, we present two sets of FSI comparison images from April 6, 2024 (near perihelion) and January 9, 2024 (near aphelion), with the same error bound of 1e-2. We find that at high compression ratio (190 \times and 860 \times), still no significant differences are visible in the perihelion image (Fig. 7 a), while the aphelion reconstructed image shows discrepancies in coronal morphology (appearing as discontinuities in the coronal structure, Fig. 7 b). This occurs because at perihelion, the closer distance provides approximately 3 times higher resolution of the solar disk compared to aphelion, resulting in lower compression ratios near perihelion. We extracted intensity curves from a circle at 1.05 solar radii in the FSI images and compared the correlation coefficients between original and compressed image intensity curves (shown in the right panel). We found that despite the apparent morphological information loss in the aphelion images (with higher compression ratios), they still maintain high correlation coefficients with the original image intensity curves. In Fig. 7 c, we present the pearson correlation coefficients between the original and reconstructed intensity distributions at 1.05 solar radii over the complete 30-month FSI dataset under three error bounds. At an error bound of 1e-3, the correlation remains extremely close to 1 (average of 0.99998). Even under a looser bound of 1e-2, although the coefficient is lower, our visual inspection confirms that the resulting reconstruction errors still remain within acceptable levels.

Unlike traditional coronal EUV imagers, FSI has an unprecedentedly large FOV: $(228')^2$ (Rochus et al. 2020; Berghmans et al. 2023), which has significant overlap with the Solar Orbiter coronagraph Metis (Antonucci et al. 2020). At perihelion, this FOV corresponds to $(4R_{\odot})^2$ such that the full solar disk is always seen, even at maximum off-pointing ($1R_{\odot}$). This FOV is significantly wider than the $(3.34R_{\odot})^2$ of EUVI (Howard et al. 2008) or the $(3.38R_{\odot})^2$ of SWAP (Seaton et al. 2013). When close to aphelion, this FOV corresponds to $(14.3R_{\odot})^2$, providing

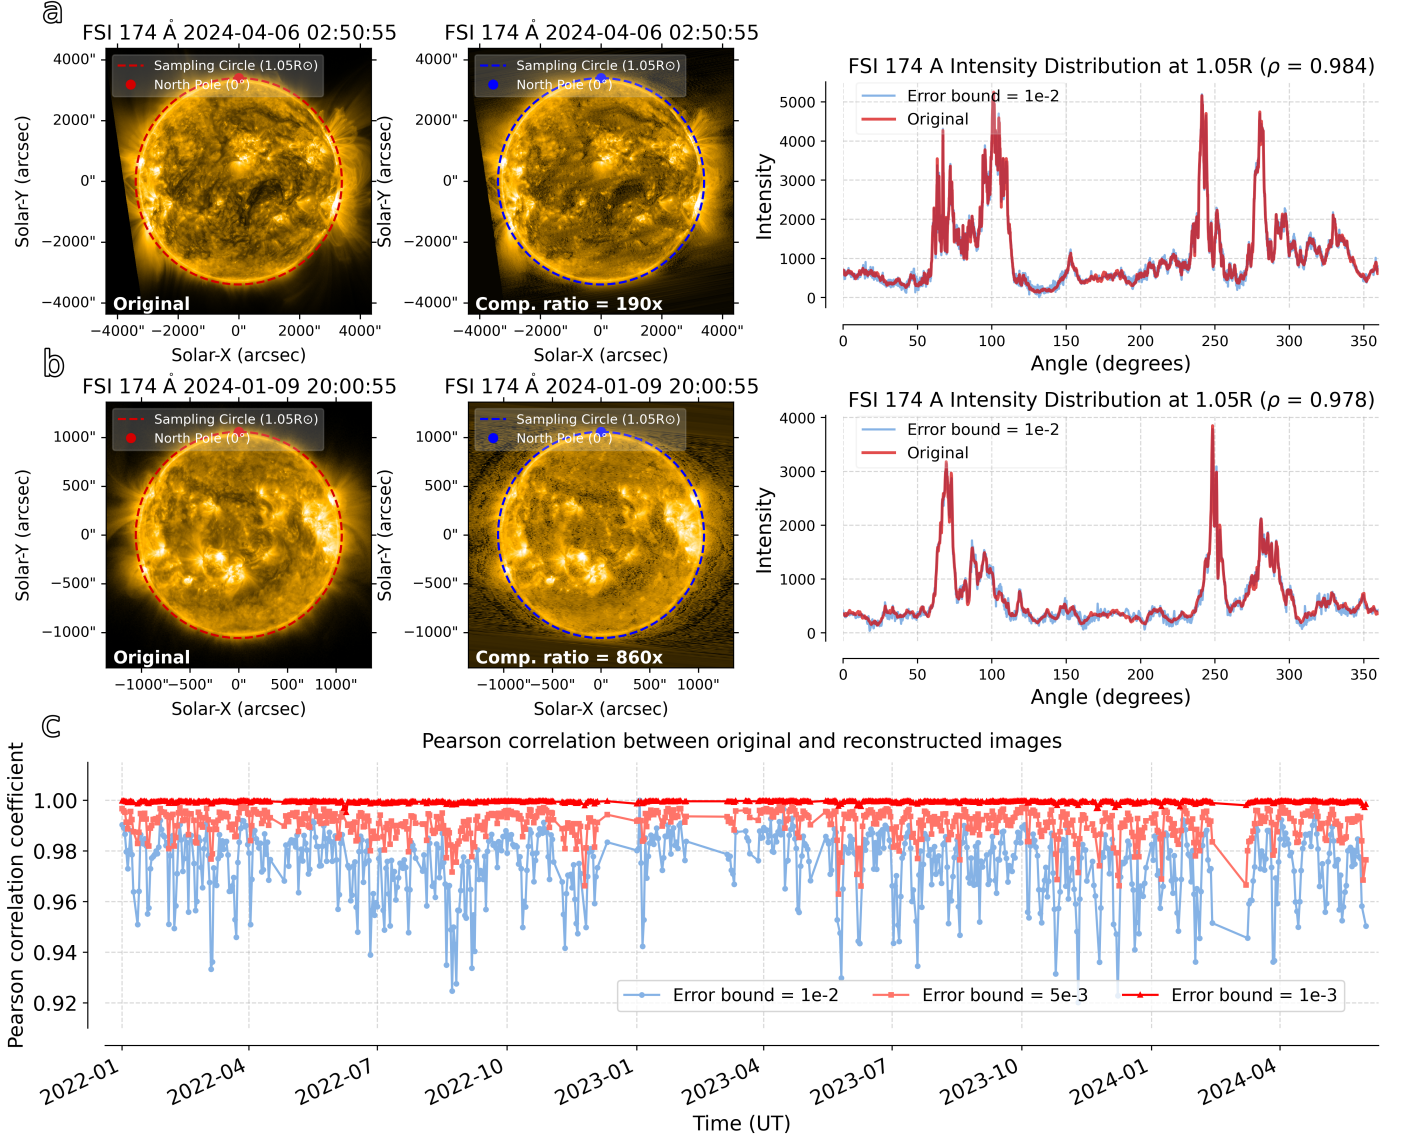


Fig. 7: Post-Hoc analysis comparison results of FSI. Panel a: Comparison image near perihelion on April 6, 2024, with red and blue dashed circles extracting intensity distributions at 1.05 solar radii from the original and reconstructed images. Results are shown in the right panel. Panel b: Comparison image near aphelion on January 9, 2024, with red and blue dashed circles extracting intensity distributions at 1.05 solar radii from the original and compressed images. Results are displayed in the right panel. Panel c: Correlation coefficients of intensity distributions at 1.05 solar radii between original and reconstructed images under three different error bounds, based on 2.5 years of FSI data.

unique opportunities to image the middle corona and eruptions that transit through this region.

Our post-analysis demonstrates that for FSI with such dynamically varying field of view and observational targets, our SolarZip algorithm achieves compression ratios of nearly 200× for perihelion FSI images at the error bound of 1e-2, while aphelion FSI images reach ultra-high compression ratios exceeding 800×. Considering the dynamic variations of FSI data, user can lower the error bound when Solar Orbiter approaches the Sun to achieve higher fidelity, and increase the error bound when it moves away from the Sun to reduce fidelity while maintaining higher compression ratios, keeping the overall compression performance within an optimal range.

5.2.2. HRI_{EUV} reconstructed image analysis

The HRI_{EUV} plate scale is 0.492'', which implies unprecedented ultra-high resolution EUV observations. On April 5, 2024, Solar Orbiter reached a distance of 0.29 AU from the Sun, giving (single) pixel values on the Sun of (105 km)² for HRI_{EUV}. This unparalleled resolution of HRI_{EUV} provided us with an opportunity to study small-scale dynamic structures in the solar corona. We selected the frame at 22:53:48 UT as a representative image to analyze whether the reconstructed image meets the actual scientific requirements (Fig. 8 a). We selected two regions for comprehensive analysis, covering two well-studied phenomena widely present in the solar atmosphere: solar jets and prominences.

Solar jets, defined as collimated, beam-like plasma ejections along magnetic field lines, are ubiquitous in all regions of the solar atmosphere, including active regions, coronal holes, and




Fig. 8: Post-Hoc analysis comparison results of HRI_{EUV} . Panel a: The selected demonstration image showing two representative features: a jet (area of 150 square pixels) and a prominence (area of 300 square pixels). Panel b: Comparison between the original image and images at four different compression ratios in the jet region, with annotations showing compression ratios and the sum of pixel intensities within the region. Panel c: Comparison between the original image and images at four different compression ratios in the prominence region, with annotations showing compression ratios and the sum of pixel intensities within the region.




Fig. 9: Panels a-h: Difference maps between original and compressed images, labeled with corresponding panels from Fig. 8. Right side of difference maps: Pixel difference distribution of the difference maps, annotated with median and standard deviation. Black solid line indicates the median position, blue dashed lines represent the range of one standard deviation on either side of the median.

quiet-Sun regions (Raouafi et al. 2016; Shen 2021; Joshi 2021; Tan et al. 2022, 2023; Sterling et al. 2024). Prominences, as widely present magnetized plasma structures in the solar atmosphere, exhibit rich dynamic characteristics (Chen 2011; Webb

& Howard 2012; Warmuth 2015; Shen et al. 2020; Asai et al. 2012). We compared the reconstructed and original images of the jet region and prominence region under different eb values (corresponding to different compression ratios, Fig. 8 b and c).

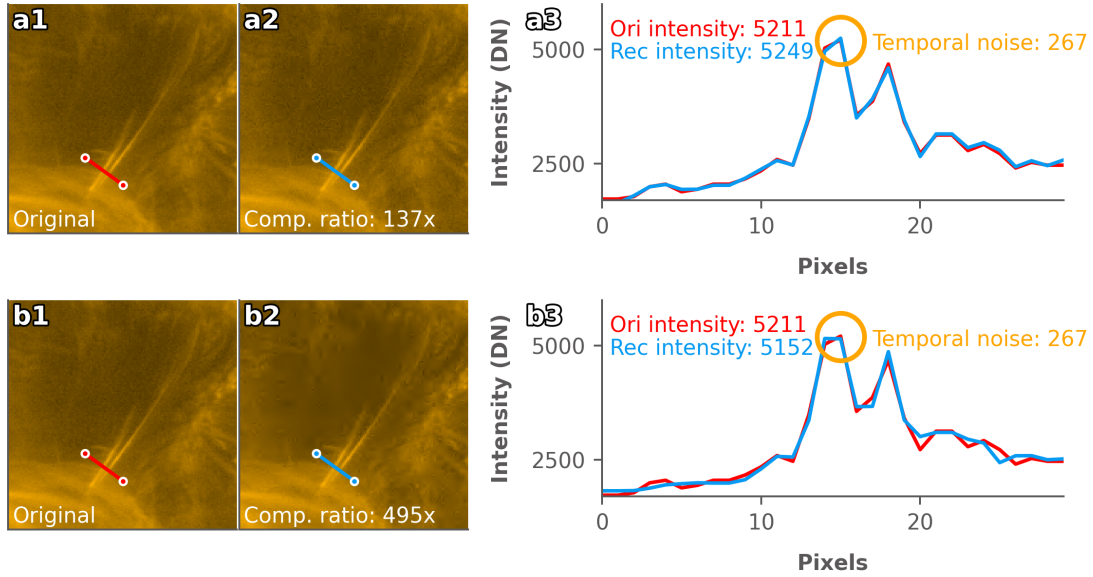


Fig. 10: Comparison of algorithm-induced errors versus HRI_{EUV} temporal noise errors. Panel a: Intensity distributions extracted across the jet structure from both the original image and the image with 137 \times compression ratio. Panel b: Intensity distributions extracted across the jet structure from both the original image and the image with 495 \times compression ratio. The intensity maxima in the cross-jet intensity distributions are marked for each image.

Firstly we can find that the intensity and of both the original image and the reconstructed image at different compression rates remain highly consistent. For the smaller jet region (150 square pixels), we observed noticeable blurring effects beginning at a compression ratio of 831 \times , while both 137 \times and 495 \times compression ratios maintained morphological consistency. For the larger prominence region (300 square pixels), our morphology-based visual inspection show that the reconstructed images maintained excellent prominence structural features compared to the original image, even at the highest compression ratio of 1286 \times .

To further analyze differences between original and reconstructed HRI images, we employed two complementary methods. First, we generated difference maps by subtracting the reconstructed images from the originals, highlighting structural discrepancies in selected regions containing solar jets (Fig. 9 a-d, left panels) and prominences (Fig. 9 e-h, left panels). This enabled direct morphological comparison of fine-scale features. Second, we computed the normalized intensity difference (1-(original/reconstructed)) and plotting its histogram (Fig. 9, right panels). We computed the median and standard deviation to characterize the compression artifacts, with vertical lines marking the median and $\pm 1\sigma$ boundaries. The two structures exhibited different sensitivities to different compression ratios. At compression ratios of 137 \times and 495 \times , difference images of the jet region showed no distinct structural features, appearing primarily as random noise. However, at 1286 \times , elongated jet structures emerged in the difference map (Fig. 9 d), indicating significant discrepancies at the jet edge. In contrast, the prominence region difference images only began to show subtle prominence structures at the 1286 \times (Fig. 9 h), yet the histogram at this compression ratio maintained a relatively narrow distribution, upporting our visual assessment of preserved morphological features.

In our analysis of the jet region, we found that the difference image at a compression ratio of 137 \times showed almost no jet structures (Fig. 9 a), instead exhibiting features of noise signals. We therefore sought to compare the algorithm-induced errors of the jet structures with the inherent temporal noise of the HRI_{EUV}

instrument itself. The temporal noise on repeated HRI_{EUV} pixel values is typically dominated by the sensor read noise and the photon shot noise according to the formula (Kraaikamp et al. 2023):

$$s^2 = r^2 + I * t * a / n \quad (7)$$

where s is the uncertainty on measured value, r is the read-out noise (2 DN), I is the measured intensity (DN/s), t is the exposure time (s), a is the photon to DN conversion factor (6.85 DN/photon), and n is the sample size (number of pixels over which the intensity is averaged). As shown in Fig. 10, we compared the algorithm-induced errors at compression ratios of 137 \times and 495 \times with the temporal noise of the HRI_{EUV} instrument. For the intensity distribution across the jet shown, the peak intensities in the reconstructed and original images were 5249:5211 and 5152:5211, indicating errors of 38 DN/s (137 \times) and -59 DN (495 \times), respectively. For $I = 5211$ DN/s with an exposure time of $t = 2.0$ s, the uncertainty on the measured value is 267 DN. This demonstrates that even at a compression ratio of 495 \times , the error in the peak intensity of the jet in the reconstructed image due to the algorithm remains significantly smaller than the measurement uncertainty caused by the HRI_{EUV} temporal noise.

Our comprehensive analysis indicates that for this set of ultra-high resolution HRI_{EUV} observations near perihelion, SolarZip can effectively achieve compression ratios of several hundred times without affecting scientific analysis. Specifically, for high-contrast dynamic structures represented by solar jets, compression ratios of around 500 \times can be achieved (corresp. 2e-2 error bound). For quiescent prominences, even higher compression ratios of approximately 800 \times are possible (corresp. 3e-2 error bound).

6. Conclusions

This paper presents SolarZip, an efficient and adaptive compression and evaluation framework specifically designed for solar

EUV data from solar missions. To our knowledge, this is the first study to apply and systematically analysis advanced error-bounded lossy compressors (SZ, ZFP, and SPERR) on solar EUV data, demonstrates their significant advantages over traditional methods. However, their inherent limitations motivates us to design an adaptive hybrid compression strategy, dynamically selecting optimal decorrelation model based on observational scenarios. Then we further optimize the interpolation predictor to enhance compression quality. SolarZip achieving unprecedented compression ratios of up to 800× for FSI data and 500× for HRI_{EUV} data—surpassing traditional algorithms by 3–50×. Our two-stage evaluation framework ensures that compressed data remains suitable for critical scientific research by integrating strict error control with downstream scientific workflows. The simulation experiments based on Solar Orbiter hardware conditions confirm that SolarZip can reduce data transmission time by a factor of 270, addressing a critical bottleneck in deep space solar missions. The source code will be made publicly available at: [SolarZip](#).

The test data used in this study are derived from publicly available data products that have already undergone initial compression and decompression, which may affect the evaluation of compression performance. We hope to conduct future tests on more original data, as experience suggests that compression performance are typically better. Future work will focus on algorithmic enhancements tailored to the specific properties of solar data. For example, we aim to introduce a Region of Interest (ROI) approach, enabling the application of tighter error constraints in scientifically critical regions to achieve superior compression performance. Moreover, the versatility of our compression framework suggests its potential applicability to other types of astronomical datasets, such as radio observations, which represents a promising avenue for future exploration.

Acknowledgements. Solar Orbiter is a space mission of international collaboration between ESA and NASA, operated by ESA. The EUI instrument was built by CSL, IAS, MPS, MSSL/UCL, PMOD/WRC, ROB, LCF/IO with funding from the Belgian Federal Science Policy Office (BELPSO); the Centre National d'Etudes Spatiales (CNES); the UK Space Agency (UKSA); the Bundesministerium für Wirtschaft und Energie (BMWi) through the Deutsches Zentrum für Luft- und Raumfahrt (DLR); and the Swiss Space Office (SSO). D.T. would like to acknowledge support from the National Natural Science Foundation of China (Grant Nos. 62032023, and T2125013) and the Innovation Funding of ICT, CAS (Grant No. E461050). B.Z. was supported by the NSFC Fund (042274216). The AIP team was supported by the German Space Agency (DLR), grant number 50 OT 2304. This research used the SunPy ([The SunPy Community et al. 2020](#); [Mumford et al. 2020](#)) and [NicePlots](#) software package to present the observation results.

References

- Ainsworth, M., Tugluk, O., Whitney, B., & Klasky, S. 2019, *SIAM Journal on Scientific Computing*, 41, A2146
- Antonucci, E., Romoli, M., Andretta, V., et al. 2020, *Astronomy & Astrophysics*, 642, A10
- Asai, A., Ishii, T. T., Isobe, H., et al. 2012, *ApJ*, 745, L18
- Baker, A. H., Hammerling, D. M., Mickelson, S. A., et al. 2016, *Geoscientific Model Development*, 9, 4381
- Baker, A. H., Hammerling, D. M., & Turton, T. L. 2019, in *Computer Graphics Forum*, Vol. 38, Wiley Online Library, 517–528
- Baker, A. H., Xu, H., Dennis, J. M., et al. 2014, in *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*, 203–214
- Baker, A. H., Xu, H., Hammerling, D. M., Li, S., & Clyne, J. P. 2017, in *High Performance Computing: ISC High Performance 2017 International Workshops, DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG, P³MA, VHP, Visualization at Scale, WOPSSS, Frankfurt, Germany, June 18–22, 2017, Revised Selected Papers 32*, Springer, 30–42
- Berger, T. 2003, *Wiley Encyclopedia of Telecommunications*
- Berghmans, D., Antolin, P., Auchère, F., et al. 2023, *Astronomy & Astrophysics*, 675, A110
- Cappello, F., Di, S., Li, S., et al. 2019, *The International Journal of High Performance Computing Applications*, 33, 1201
- Chege, J., Koopmans, L., Offringa, A., et al. 2024, *Astronomy & Astrophysics*, 692, A211
- Chen, P. F. 2011, *Living Reviews in Solar Physics*, 8, 1
- Di, S. & Cappello, F. 2016, in *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, IEEE, 730–739
- Di, S., Liu, J., Zhao, K., et al. 2024, *arXiv preprint arXiv:2404.02840*
- Diffenderfer, J., Fox, A. L., Hittinger, J. A., Sanders, G., & Lindstrom, P. G. 2019, *SIAM Journal on Scientific Computing*, 41, A1867
- Fischer, C. E., Müller, D., & De Moortel, I. 2017, *Solar Physics*, 292, 16
- Howard, R. A., Moses, J., Vourlidas, A., et al. 2008, *Space Science Reviews*, 136, 67
- Jin, S., Grosset, P., Biver, C. M., et al. 2020, in *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, IEEE, 105–115
- Jin, S., Pulido, J., Grosset, P., et al. 2021, in *Proceedings of the 30th International Symposium on High-Performance Parallel and Distributed Computing*, 45–56
- Joshi, R. 2021, PhD thesis, Kumaun University, India
- Kraaikamp, E., Gissot, S., Stegen, K., et al. 2023, *Solo/EUI Data Release 6.0 2023-01*, <https://doi.org/10.24414/z818-4163>, published by Royal Observatory of Belgium (ROB)
- Leung, Y. K. & Apperley, M. D. 1994, *ACM Transactions on Computer-Human Interaction (TOCHI)*, 1, 126
- Li, S., Lindstrom, P., & Clyne, J. 2023, in *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, IEEE, 1007–1017
- Liang, X., Di, S., Li, S., et al. 2019, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 1–26
- Liang, X., Di, S., Tao, D., Chen, Z., & Cappello, F. 2018a, in *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, IEEE, 179–189
- Liang, X., Di, S., Tao, D., et al. 2018b, in *2018 IEEE International Conference on Big Data (Big Data)*, IEEE, 438–447
- Liang, X., Zhao, K., Di, S., et al. 2022, *IEEE Transactions on Big Data*, 9, 485
- Lindstrom, P. 2014, *IEEE transactions on visualization and computer graphics*, 20, 2674
- Liu, J., Di, S., Zhao, K., et al. 2022, in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, 1–15
- Liu, X., Liu, Y., Yang, L., et al. 2024a, *Publications of the Astronomical Society of the Pacific*, 136, 075001
- Liu, Z., Jiang, P., Zeng, F., Bian, H., & Toe, T. T. 2024b, in *2024 4th International Conference on Computer Communication and Artificial Intelligence (CCAI)*, IEEE, 58–64
- Marinrodrga, C. G., Pacros, A., Strandmoe, S., et al. 2021, *Astronomy & Astrophysics*, 646, A121
- Müller, D., St. Cyr, O. C., Zouganelis, I., et al. 2020, *A&A*, 642, A1
- Mumford, S., Freij, N., Christe, S., et al. 2020, *Journal of Open Source Software*, 5, 1832
- Patel, H., Itwala, U., Rana, R., & Dangarwala, K. 2015, *International Journal of Engineering Research and Technology*, 4, 926
- Peters, S. M. & Kitaef, V. V. 2014, *Astronomy and Computing*, 6, 41
- Pulido, J., Lukic, Z., Thorman, P., et al. 2019, in *Journal of Physics: Conference Series*, Vol. 1290, IOP Publishing, 012008
- Raouafi, N. E., Patsourakos, S., Parlat, E., et al. 2016, *Space Sci. Rev.*, 201, 1
- Rice, R. & Plaunt, J. 1971, *IEEE Transactions on Communication Technology*, 19, 889
- Rochus, P., Auchère, F., Berghmans, D., et al. 2020, *Astronomy & Astrophysics*, 642, A8
- Seaton, D., Berghmans, D., Nicula, B., et al. 2013, *Solar Physics*, 286, 43
- Shen, Y. 2021, *Proceedings of the Royal Society of London Series A*, 477, 217
- Shen, Y., Li, B., Chen, P., Zhou, X., & Liu, Y. 2020, *Chinese Science Bulletin*, 65, 3909
- Sterling, A. C., Panesar, N. K., & Moore, R. L. 2024, *ApJ*, 963, 4
- Tan, S., Shen, Y., Zhou, X., et al. 2022, *MNRAS*, 516, L12
- Tan, S., Shen, Y., Zhou, X., et al. 2023, *MNRAS*, 520, 3080
- Tao, D., Di, S., Chen, Z., & Cappello, F. 2017, in *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, IEEE, 1129–1139
- Tao, D., Di, S., Liang, X., Chen, Z., & Cappello, F. 2019, *IEEE Transactions on Parallel and Distributed Systems*, 30, 1857
- Taubman, D. S. & Marcellin, M. W. 2002, *Proceedings of the IEEE*, 90, 1336
- The SunPy Community, Barnes, W. T., Bobra, M. G., et al. 2020, *The Astrophysical Journal*, 890, 68
- Vohl, D., Fluke, C. J., & Vernardos, G. 2015, *Astronomy and Computing*, 12, 200
- Wallace, G. K. 1991, *Communications of the ACM*, 34, 30
- Warmuth, A. 2015, *Living Reviews in Solar Physics*, 12, 3
- Webb, D. F. & Howard, T. A. 2012, *Living Reviews in Solar Physics*, 9, 3

- Xie, H., West, R. A., Seignovet, B., et al. 2021, *Journal of astronomical telescopes, instruments, and systems*, 7, 028002
- Zafari, A., Khoshkhahtinat, A., Grajeda, J. A., et al. 2023, *IEEE Transactions on Aerospace and Electronic Systems*, 60, 918
- Zafari, A., Khoshkhahtinat, A., Mehta, P. M., et al. 2022, in *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, IEEE, 198–205
- Zhang, B., Guo, L., Tian, J., et al. 2025, in *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, 557–559
- Zhao, K., Di, S., Dmitriev, M., et al. 2021, in *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, IEEE, 1643–1654
- Zhao, K., Di, S., Liang, X., et al. 2020, in *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing*, 89–100

Appendix A: Advanced lossy compressor

Appendix A provides a detailed description of the four advanced error-bounded lossy compression algorithms employed for evaluation and comparison in this study. We first present their workflows and then highlight their respective characteristics.

A.1. SZ2

SZ2 (Liang et al. 2018a) refers to the second-generation SZ compression algorithm. Which is a prediction-based error-bounded lossy compressor.

Workflow: SZ2's compression has four main steps. Firstly, it divides raw data to be compressed into small blocks. For each of these blocks, a separate prediction function is generated. Secondly, SZ2 quantizes the data with specified error bound. Thirdly, employing Huffman encoding to encode the quantization index. Finally, lossless compression methods are used to further improve compression ratio.

Insight: SZ2 compression process is simple and efficient, achieving notable compression performance and speed on most scientific datasets. However, due to its block-wise prediction approach, its effectiveness is limited for high-dimensional and non-linear data.

A.2. SZ3

SZ3 (Liang et al. 2022; Zhao et al. 2021) uses a modular approach to compress the data. In fact, it is not only a compression software but also a flexible framework allowing users to customize specific compression pipelines according to their datasets or use cases.

Workflow: SZ3 compression pipeline is composed of five stages: preprocessing, prediction, quantization, variablelength encoding, and lossless compression. Preprocessing stage start by transforming and shaping the raw data to where is more easily to compressed. The second stage is prediction. For different domain datasets, SZ developers have developed many predictors, including Lorenzo, linear regression, and dynamic spline interpolation. Thirdly, the error produced by the predictor is quantized. Fourthly, the quantized error data is encoded, sHRI_{EUV} linking their size. Fifthly, the now encoded data is losslessly compressed, reducing the size further more.

Insight: With an advanced predictive model and a flexible modular design, SZ3 significantly enhances compression performance and adaptability. However, these improvements come with relatively higher computational costs.

A.3. ZFP

ZFP (Lindstrom 2014; Diffenderfer et al. 2019)) is a transform-based error-bounded lossy compressor.

Workflow: ZFP splits the whole dataset into many fixed-size blocks (e.g., 4×4×4 for a 3D dataset) These blocks are then individually compressed. ZFP compressor executes four steps in each block. The first step align the value in block to a common exponent and convert the floating-point values to a fixed-point representation. The next step uses orthogonal block transform to decorrelate data. Thirdly, it orders the transform coefficient by expected magnitude. Finally, it encodes the coefficients to reduce data size.

Insight: In general, ZFP features high compression and decompression performance because of the performance optimization

strategies in its implementation. However, as a block-wise compressor, ZFP faces the same limitations as SZ2.

A.4. SPERR

SPERR (Li et al. 2023) is a transform-based lossy compressor based on the CDF9/7 discrete wavelet transform and SPECK encoding algorithm.

Workflow: The compression pipeline of SPERR includes four stages: (1)CDF9/7 wavelet transform; (2) SPECK lossy encoding of wavelet coefficients; (3) outlier encoding (only in error-bounding mode); (4) zstd postprocessing of compressed data (optional).

Insight: The advantage of SPERR is that the hierarchical multi-dimension DWT in SPERR can effectively capture the relevance between data points, which brings a high compression ratio after the SPECK encoding. One limitation of SPERR is that the wavelet transform and the SPECK encoding processes have high computational costs, and hence its (sequential) execution speed is relatively low, typically around 30% of SZ3.

Appendix B: Anchor points interpolations with error bound auto-tuning.

A issue is that classical spline interpolation applies the same error bound to all prediction levels, which does not fully account for their relative importance. As shown in the fig. B.1, the predicted data points in Level 1 participate in five subsequent prediction steps. Consequently, earlier-predicted data points should be considered more important.

First, we apply different interpolation methods at different levels of interpolation prediction. Specifically, the interpolation types include linear interpolation and cubic interpolation. In our two-dimensional data, the interpolation process is essentially carried out through multiple one-dimensional interpolation operations. Since there are two dimensions (dim0 and dim1), two distinct interpolation orderings are possible, and we select the optimal interpolation sequence.

Next, we set different error bounds at different levels of interpolation (as opposed to the unified error bounds used in SZ3). In our two-dimensional data, 75% of the data points fall within the lowest interpolation level (Level 1), which are predicted by higher-level reconstructed data points, while the remaining 25% of the data points are predicted at higher levels. Therefore, setting smaller error bounds at higher levels helps ensure overall prediction accuracy, thus improving compression quality. The level-wise error bounds e_l are dynamically adjusted based on equation 6. The parameters α and β are introduced, where e represents the global error bound set by the user. We perform offline testing with parameter sets $\alpha = \{1, 1.5, 2\}$ and $\beta = \{2, 3, 4\}$, comparing the bit-rate and PSNR values across different parameter configurations. Ultimately, we select $\alpha = 1.5$ and $\beta = 4$ as our optimal parameters.

Appendix C: Data availability

Solar Orbiter data is publicly available through the Solar Orbiter Archive (<https://soar.esac.esa.int/soar/>).

The sample data used to generate the figures in this work are publicly available at the following link: [SolarZip-TestData](#). The dataset includes the raw data and the decompressed reconstructed data for each figure. For instance, in Fig. 8 and 9, we performed a case study using an HRI_{EUV} observation from April

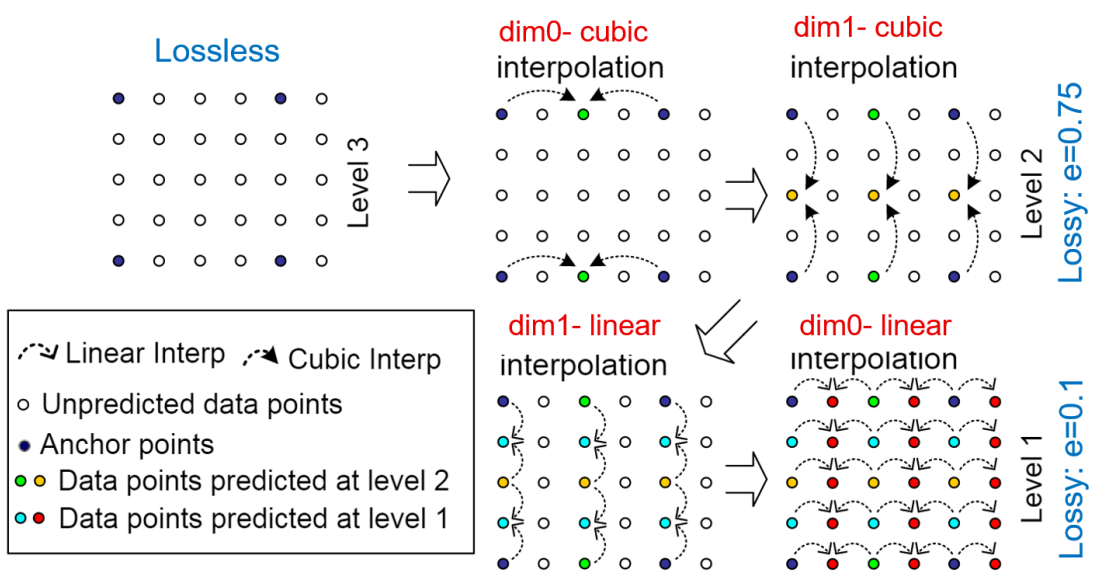


Fig. B.1: The level-wise anchor points based dynamic spline interpolation. Using a two-dimensional dataset as an example, the three stages of the interpolation-prediction algorithm are illustrated. Starting from Level 3, interpolation is performed based on the anchor data, where the retained anchor points divide the data into fixed-size blocks. In two-dimensional data, there are two interpolation directions, and each interpolation step can utilize either a linear or cubic interpolation function. The algorithm automatically optimizes the selection process. It is evident that at the more critical Level 2, a smaller error bound is assigned, which further enhances the accuracy of the interpolation prediction.

5, 2024. This dataset contains the original image as well as four reconstructed images with compression error bounds of 1×10^{-2} , 2×10^{-2} , 3×10^{-2} , and 5×10^{-2} .

Researchers interested in the data or the SolarZip algorithm itself are encouraged to contact the authors for further information.