



psit 1.0: a system to compress Lagrangian flows

Alexander Pietak¹, Langwen Huang², Luigi Fusco^{2,3}, Michael Sprenger^{4,★}, Sebastian Schemm^{4,★}, and Torsten Hoefler²

¹Department of Mathematics, ETH Zürich, Zurich, Switzerland

²Scalable Parallel Computing Lab, ETH Zürich, Zurich, Switzerland

³Microsoft Ibérica, Madrid, Spain

⁴Department of Environmental Systems Science, ETH Zürich, Zurich, Switzerland

★These authors contributed equally to this work.

Correspondence: Alexander Pietak (apietak@ethz.ch)

Received: 19 February 2025 – Discussion started: 14 March 2025

Revised: 10 February 2026 – Accepted: 26 March 2026 – Published: 13 May 2026

Abstract. Meteorological simulations produce large amounts of data, which can represent a challenge when trying to store, share, and analyze it. As weather and climate models increasingly simulate the atmosphere at higher spatio-temporal resolution, it becomes imperative to compress the data effectively. While compression algorithms exist for weather data stored in a gridded Eulerian frame, there are, to date, no specialized alternatives for data stored in the Lagrangian frame. In this study, we present psit, a system to compress weather data stored in the Lagrangian frame. The system works by mapping the trajectories to a grid structure, performing additional encodings on these, and passing them to either the JPEG 2000 image compression algorithm or SZ3. The specialty of the algorithm is the mapping phase and the following encodings, which generate the grids in a way that allows the aforementioned compression algorithms to perform well. To gauge the performance of psit, we test a variety of metrics. We demonstrate that in the majority of cases, equivalent or superior compression performance is attained through the utilization of psit as opposed to naive compression with ZFP or SZ3. We also compare compression with measurement inaccuracies. Here, we show that the density of 168 hour long trajectories compressed with a ratio in the range of 30 to 40 behaves similarly to trajectories calculated from uncompressed wind fields with additional random perturbations with magnitude of 0.1 m s^{-1} in the horizontal and around $6 \times 10^{-3} \text{ Pa s}^{-1}$ in the vertical component. Additionally, we conduct two case studies in which we discuss the impact of compression on the study of warm conveyor belts associated with extratropical

cyclones and the impact of compression on the radioactive plume prediction of the Fukushima incident in 2011.

1 Introduction

With advancements in parallel computing and the introduction of larger and more powerful HPC systems, the size and complexity of today's meteorological models have increased significantly. Such an increase in complexity leads to a rise in the amount of output data, reaching the range of dozens of petabytes (Hoefler et al., 2024, 2023). Storing and analysing such a massive amount of data poses significant challenges. Consequently, different approaches have been developed to address this issue. One approach works by reducing the number of elements that are stored (Tintó Prims et al., 2024). In the case of weather data, this can be achieved, for example, by reducing the grid resolution of the stored data or decreasing the temporal interval at which data is stored. This approach is already in use for the ERA5 weather dataset (Tintó Prims et al., 2024). A second approach that is commonly used focuses not on reducing the number of elements but on decreasing the space occupied by each of them. This is the field of compression algorithms, which reduce data size by exploiting structure present in the underlying data, a technique commonly used in multimedia, as seen in formats like JPEG (Wallace, 1991) and H.265 (Sullivan et al., 2012). For the compression of scientific data, we usually prefer lossy compression, as, due to the inherent noise present,

lossless compression algorithms will not be able to achieve large compression ratios (in the range of 2 to 4 times). We will therefore focus on lossy compression techniques.

Several lossy compression algorithms for the compression of multidimensional floating-point data have been developed. A (non exhaustive) list includes: SZ3 (Liang et al., 2023b, 2018; Zhao et al., 2021), ZFP (Lindstrom, 2014; Diefenderfer et al., 2019), and THRESH (Ballester-Ripoll et al., 2019; Ballester-Ripoll and Pajarola, 2015). Previous research by Tintó Prims et al. (2024) has demonstrated the effectiveness of these floating-point compression algorithms in compressing weather data in the Eulerian frame. There are also specialized compression algorithms developed specifically for weather data, such as VAEformer (Han et al., 2024), an autoencoder-based compression algorithm designed for the compression of the ERA5 weather dataset. Hence, a wide variety of compression algorithms exist that can be utilized for the compression of weather data in the Eulerian frame.

On the other hand, meteorologists not only work with the Eulerian representation but also often consider a Lagrangian one. In this representation, we do not have a multidimensional grid but instead track the position and properties, like temperature, of small air parcels carried by the wind, resulting in so called trajectories. This representation is crucial for certain types of analysis, such as identifying warm conveyor belts (Wernli and Davies, 1997b; Joos and Wernli, 2012; Schemm et al., 2013) or stratosphere-troposphere exchanges (Holton et al., 1995; Škerlak et al., 2014). However, here we also face the problem of an increase in output data and therefore seek to compress it. Compared to multidimensional grid data, the landscape of compression algorithms for weather data trajectories is much poorer. There are mainly two fields related to this, first is the compression of trajectory data, where research mostly focuses on data generated from GPS tracking of vessels or, more generally, the movement of objects in space over many time steps (Makris et al., 2021). Examples of such algorithms include Dead-Reckoning, Douglas-Peucker, and Time-Ratio (Makris et al., 2021), all of which reduce data volume by strategically removing individual time steps while preserving the general trajectories. However, these methods are not suitable for compressing trajectories originating from weather data for two main reasons. First, the number of time steps in weather trajectories is typically much lower than in GPS trajectories. For GPS trajectories, the number of time steps might be in the 10 000 range, while for weather trajectories, depending on the use case, one might work with only a few dozen (e.g., hourly to six hourly temporal resolution). Additionally, removing time steps from weather trajectory data may not be permissible due to the specific requirements of subsequent research. Secondly, weather trajectory data often includes additional variables, such as temperature or potential vorticity, that are of interest along the trajectories. The above mentioned trajectory compression algorithms do not account for these additional variables, only for the position,

which makes them unsuitable for compressing such weather data. The second research field related to this is the compression of unstructured mesh data. Different methods like (Liang et al., 2023a; Ren et al., 2024; Wu et al., 2025) exists and a recent survey has been made by Di et al. (2025). But quite a lot of them are not directly designed for time series data, but focus on other aspects such as point clouds (Quach et al., 2022) or medical data (Al-Salamee and Al-Shammary, 2021). The research that focuses on time series data (Chiarot and Silvestri, 2023) mostly is in the field of internet of things and not on weather data. So we can see that there appears to be a gap when it comes to the compression of trajectories originating from weather data.

To address this gap, we present psit, a lossy compression method for Lagrangian flow data. Psit works by transforming the trajectory data into 2D grids and passing those grids to existing compression algorithms, specifically JPEG 2000 (Skodras et al., 2001) and SZ3, where JPEG 2000 has shown its potential in internal research and is used by the visualization community due to its adaptive scalability property (Woodring et al., 2011). This approach leverages the refined performance of these compression algorithms to create a compression pipeline specifically designed for trajectory weather data.

2 Method

The general idea behind psit is to take the trajectories and intelligently transform them so that they can be passed to existing compression algorithms. This general pipeline is illustrated in Fig. 1. The input expected by the dense data compressors (JPEG 2000 and SZ3) consists of 2D arrays (i.e., grids), which ideally should be smooth, i.e., the difference between neighboring entries of the array should be small. Note that a single 2D array (representing one channel) can be passed to the JPEG 2000 algorithm, as it is able to handle an arbitrary number of channels (individual 2D arrays) and is not limited to the default three color channels (red, green, and blue) used for images. Therefore, in the first step (the mapping phase), the trajectories are transformed into such grids. During this mapping, it is crucial to assure the resulting grids smoothness, as greater smoothness allows both compression algorithms to perform more effectively. Once the mapping step is completed, encoding schemes can be applied to them, namely delta and color encoding. These encoding schemes modify the grids in a way that leads to better compression performance for the previously mentioned compression algorithms. In the next step, the encoded grids are passed to a dense data compressors to perform the compression, and the compressed data is subsequently stored to disk. These four steps form the basis of psit and are discussed in more detail in the rest of this section.

The decompression phase mirrors the compression process: the data is loaded from disk and then decompressed us-




Figure 1. The compression pipeline of psit. In the first step, the mapping phase, trajectory data is converted into grid data. After these grids have been created they can be manipulated in order to have better compression performance, namely delta and color encoding. Then as a final step the grids are compressed with either JPEG 2000 or SZ3 and stored to disk. The decompression works the same, but in the opposite direction.

ing the same algorithm employed during compression (e.g., JPEG 2000). After this, the encoded grids are decoded (delta and color decoding) to obtain the mapped 2D grids. Finally, these 2D grids are transformed back into trajectories using an inverse mapping of the original one. After this step, the decompression is complete, and the trajectory data should closely match the original data, with only the errors introduced by the compression algorithm differentiating them.

2.1 Input data

In the following, we will provide a definition of what a trajectory and a possible input file are in order to establish a common basis for later discussions. We define a trajectory as a sequence of positional variables (longitude, latitude, and pressure, which indicates the height) defining a 3D path on Earth, along with additional sequences that store variables like temperature or humidity along it. Each entry in these sequences represents one time step. Multiple such trajectories (in the range of a few million) can then be collected together to form a trajectory file, which would function as input to our compression algorithm.

For example, an input file might consist of one million trajectories spanning 13 time steps, with each trajectory also storing temperature and humidity. At the first time step, the trajectories are uniformly initialized over 26 pressure levels spanning the entire globe, with some predefined distance between the different starting points. This would result in a file storing a total of five million individual sequences (five per trajectory: 3 positional and 2 data variables), where each sequence stores 13 elements.

2.2 Mapping

The mapping phase is the first step in the compression pipeline and focuses on transforming the trajectories into grids while ensuring that the resulting grids are smooth. To achieve this, a method for storing the trajectory data as grids must be established. A simple approach is used: each trajectory is associated with at least one grid point, as visually represented in Fig. 2. The next step consists of translating the data from the trajectories into grids. For this, we first create an individual grid for each time step, and because a trajectory

consists of multiple data variables (3 positional and an arbitrary number of additional ones), we also create an individual grid for each of those. After all these grids are created, we simply set the values of the grid points to the value of the associated trajectories data variable at that specific time point. This way, a collection of 2D grids can be created that store the same information as the trajectories but are inherently 2D. For example, the previously introduced trajectory file consists of one million, 13 time step long trajectories, which also store temperature and humidity. This file would be converted into 65 individual grids, where each grid would have at least one million entries, e.g., a dimension of 1500×750 .

While creating such an initial association can be done trivially, there are constraints that make the task more difficult. The first and most crucial constraint is that the resulting grids need to be smooth. Secondly, we must ensure that every trajectory maps to at least one grid point. If this is not the case, a trajectory would have no representation in the compressed data and would therefore be lost during compression. These constraints, especially the smoothness one, prevent the use of naive mapping methods and require the use of more sophisticated techniques.

In our research, we analysed different mapping methods. The one that yields the best results is based on solving a minimal weight full bipartite matching problem, originating in graph theory (Burkard et al., 2012). A second promising approach is based on solving a linear program (LP) (Bertsimas and Tsitsiklis, 1997). Solving this LP becomes computational infeasible for large images, but we still include it for theoretical analysis.

2.2.1 Bipartite Mapping

For this mapping method (results in Fig. 3, with an example in Fig. 4), the problem is solved using concepts from graph theory, more specifically by solving a minimal weight full bipartite matching problem. Hence, following this approach, we have two sets of nodes: one called “workers” and the other called “tasks”, with weighted edges between them constraining which workers may be assigned to which tasks and what the cost of this assignment is. The goal is to match each worker to exactly one task while trying to keep the overall cost of the matchings as low as possible.




Figure 2. Visual representation of a mapping, each trajectory is assigned to some grid points based on their initial position, note that for visual clarity only some mappings are drawn. The way in which this association is done forms an integral part of our pipeline and is based on solving a minimal weight full bipartite matching problem.




Figure 3. Latitude grids with 550 hPa starting pressure created with the bipartite mapping method. Three different time steps are displayed the leftmost image is at the first time step, the second one is 12 h later and the rightmost one is 24 h later. In the leftmost image we have very good smoothness, which originates from the mapping method, over time this smoothness starts to degrade, as the trajectories start to diverge.

This minimum weight full bipartite matching problem can be easily applied to our use case. Namely, the workers become the trajectories and the tasks the grid points. The weights are then defined as the Euclidean distance between the trajectories and the grid points at the first time step after both are translated into 3D Cartesian space. To transform the grid points into Cartesian space, a projection onto a 3D sphere is used. This projection method should lead to the grid points being homogeneously distributed over the globe, making a simple longitude-latitude mapping unsuitable. We chose the projection presented in Sect. 4 of the paper by Calhoun et al. (2008). The trajectories are transformed into Cartesian space by taking their spherical coordinates of longitude, latitude, and pressure and converting them to Cartesian x , y , and z coordinates. This transformation leads to the creation of a spherical shell, representing a mismatch in dimensionality between the grid points and the trajectories. Therefore, the spherical shell is binned into discrete levels over its pressure, resulting in the creation of multiple spheres for each of the bins. This means that, in addition to a grid being created for each time step and data variable, individual grids are also created for each of these pressure levels. If the algorithm then minimizes the Euclidean distance between the trajectories and the grid points, the resulting grids should become smooth, as trajectories with similar starting positions should be mapped to grid points that are close to one another, and we have that trajectories starting at similar spatial positions behave rather similarly over time (hence locating them next to each other results in a smoother grid). So, by representing our mapping problem as such a graph problem and solving it, we should end up with a valid and smooth mapping.

One problem with this mapping via minimal weight full bipartite matching is that if the number of trajectories is less than the number of grid points some grid points will not be mapped, leading to holes in the resulting grid. While it theoretically is possible to have the same number of trajectories as grid points, tests showed that this can become computationally infeasible and that in practice we need more grid points than trajectories. Therefore, we need to devise a way to handle this. The hole filling works by taking all the grid points that do not get mapped and mapping their closest trajectory (Euclidean distance in Cartesian space) to them (note that this means that a trajectory may map to multiple grid points at the same time). By using this method, the holes can be filled in an easy manner that does not require any costly computations.

2.2.2 LP Mapping

We also considered using linear programming. While this approach has high theoretical appeal, its practical usability is limited by the fact that the resulting LP becomes very large and requires a lot of resources to solve, while delivering nearly identical results to the approach based on minimal weight full bipartite matching. For our use case, we will model the mapping problem as an integer linear program (ILP), for which we then show that its LP formulation is integral, i.e., solving the LP formulation instead of the ILP one delivers the same results. This allows us to use a much cheaper LP solver, resulting in a large decrease in computational resources. In order to define an integer linear program, we need to specify an objective and a set of constraints that

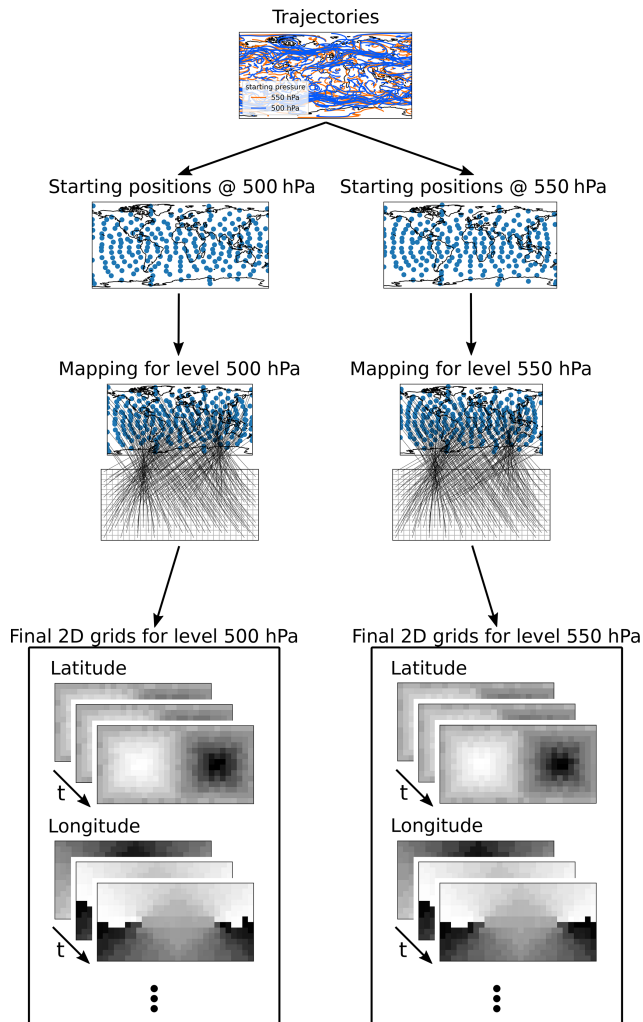


Figure 4. Example of how the bipartite mapping works. For this 464 trajectories starting at two different pressure levels are considered. In a first step we extract the starting positions of all the trajectories individually for the two different pressure levels. We then take these starting positions and create a mapping to a grid such that the resulting grids are smooth. This mapping is then used to generate a collection over time for each of the data variables which are present on our trajectories.

create a valid and smooth mapping from trajectories to grid points.

We first start with an intuitive definition of the objective and the constraints. The objective function is designed to minimize the sum of distances between the trajectories and their associated grid points. For the constraints, we say that each trajectory must map to at least one grid point and that each grid point must be mapped to. By the property of spatial continuity, a minimization of the objective function should lead to a smooth grid, while the constraints ensure that the mapping is valid. The rigorous mathematical definition is then given by:




Figure 5. Top: Longitude grids with 550 hPa starting pressure displayed over different time steps (0, 12, 24 h). At the date line there is a discontinuity which starts to warp over time. This discontinuity leads to worsened compression performance and we can use color encoding to fix this. Bottom: HSV (left) and XYZ (right) color encoded longitude grids with 550 hPa starting pressure displayed after 24 h. The date line discontinuity is gone.

Let the binary variables $x_{t,p} \in \{0, 1\}$ denote whether a trajectory $t \in T$ is mapped to grid point $p \in P$, 0 if not mapped, 1 if mapped. And let the variable $d_{t,p} \in \mathbb{R}_{\geq 0}$ denote the Euclidean distance in 3D Cartesian space (defined analogue to the bipartite mapping method) between a trajectory $t \in T$ and a grid point $p \in P$. The resulting ILP is then given in Eq. (1):

$$\text{minimize } \sum_{(t,p) \in T \times P} d_{t,p} x_{t,p} \tag{1}$$

$$\text{s.t. } \sum_{t \in T} x_{t,p} = 1 \quad \forall p \in P$$

$$\sum_{p \in P} x_{t,p} \geq 1 \quad \forall t \in T \tag{2}$$

While the variables $x_{t,p}$ are defined to be binary, making the problem an *integer LP*, we can prove (see Appendix A) that the LP formulation, achieved by defining $x_{t,p} \in [0, 1]$, is integral. This means that instead of an ILP solver a much cheaper LP solver can be used during calculation.

2.3 Color Encoding

To motivate color encoding, we must examine the longitude grids created by the previously explored mappings (Fig. 5 top), which mapping we look at (LP or bipartite) is not important as they deliver very similar results. As we can see from Fig. 5 top, there is a discontinuity line around the position of the dateline at longitude 180°, which begins to warp as time progresses. This discontinuity arises because trajectories with a longitude value of -180 (black pixels) are adjacent to trajectories with longitude values of 180 (white pixels). Such jumps in pixel values lead to worsened compression performance, as shown by experiments where we observed a decrease in RMSE error in the longitude variable by a factor of 1.4 when compressing with color encoding compared to no color encoding (using a compression factor of 15).

The way we eliminate the discontinuity line is by using multiple variables instead of one, which maps the




Figure 6. Qualitative demonstration of the discontinuity line at the date line. On the left only a singular variable (black to white) is used to represent the longitude, a discontinuity is created. On the right two variables and a cosine sine mapping is used, displayed as red and green. Using two variables it is possible to create a continuous color cycle with no discontinuity.

$[-180, 180]$ longitude range to the black to white pixels. Hence, we map the $[-180, 180]$ to a multidimensional space. If we use two or three variables, this multidimensional space can be represented by a color bar. This concept is illustrated in Fig. 6, where on the left, we have a single variable mapping the longitude from black to white, generating a discontinuity in the color wheel. On the right, we use two variables (red and green) with a sine-cosine mapping to create a color bar, which results in a continuous color wheel. This simple concept can define a variety of different color encoding methods. In our research, we explored several such methods, and the ones that provided the best performance are based on utilizing three variables, covering the entire RGB range. We call them HSV and XYZ color encoding and they are explained in the following sections.

2.3.1 HSV color encoding

The HSV color encoding method (Fig. 5 bottom left) uses the HSV (Smith, 1978) color representation, which, like RGB, is a way to describe color. However, instead of using values for red, green, and blue, it uses hue, saturation, and value to represent color. For our use case, we set the saturation and value to 1 and only vary the hue. For constant saturation and value, a varying hue generates a continuous color wheel. Another reason for setting saturation and value to 1 is that the JPEG 2000 algorithm might handle dark or washed-out colors differently, as the human eye cannot distinguish them as precisely as strong, bright colors.

The function $f(x)$ for converting a longitude value in the normalized range $[0, 1]$ to an RGB value using an HSV representation is given in Eq. (3). The inverse mapping $f_{inv}([r, g, b]^T)$, which maps an RGB color value to a hue angle, is provided in Eq. (4).

$$f(x) = \begin{bmatrix} h(5) \\ h(3) \\ h(1) \end{bmatrix}$$

with

$$h(l) = 1 - \max\{0, \min\{k, 4 - k, 1\}\}$$

$$k = (l + 6x) \bmod 6 \tag{3}$$

$$f_{inv}([r, g, b]^T) = \begin{cases} 0, & \text{if } c = 0 \\ \frac{1}{6} \cdot \left(\frac{g-b}{c} \bmod 6\right), & \text{if } v = r \\ \frac{1}{6} \cdot \left(\frac{b-r}{c} + 2\right), & \text{if } v = g \\ \frac{1}{6} \cdot \left(\frac{r-g}{c} + 4\right), & \text{if } v = b \end{cases}$$

$$v = \max\{r, g, b\}$$

$$k = \min\{r, g, b\}$$

$$c = v - k \tag{4}$$

2.3.2 XYZ color encoding

The XYZ color encoding (Fig. 5 bottom right) method considers both the longitude and latitude data variables, combining them into a single three-channel color grid. To this end, the longitude and latitude values are taken and converted into Cartesian coordinates, from which the r, g, b color channels are created from the x, y, z positional coordinates.

The function $f([x_{lon}, x_{lat}]^T)$, which maps both longitude and latitude in the normalized range $[0, 1]$ to RGB values, is given in Eq. (5). The inverse function $f_{inv}([r, g, b]^T)$, which converts RGB values back to longitude and latitude, is given by Eq. (6).

$$f([x_{lon}, x_{lat}]^T) = \frac{1}{2} \left(1 + \begin{bmatrix} \cos \hat{x}_{lat} \cos \hat{x}_{lon} \\ \cos \hat{x}_{lat} \sin \hat{x}_{lon} \\ \sin \hat{x}_{lat} \end{bmatrix} \right)$$

$$\text{where } \hat{x}_{lat} = \pi(x_{lat} - \frac{1}{2}), \hat{x}_{lon} = 2\pi(x_{lon} - \frac{1}{2}) \tag{5}$$

$$f_{inv}([r, g, b]^T) = \begin{bmatrix} \frac{1}{2\pi}(\arctan2(\hat{g}, \hat{r}) + \pi) \\ \frac{1}{\pi}(\arctan2(\hat{b}, \sqrt{\hat{r}^2 + \hat{g}^2}) + \frac{\pi}{2}) \end{bmatrix}$$

$$\text{where } \hat{r} = 2r - 1, \hat{g} = 2g - 1, \hat{b} = 2b - 1. \tag{6}$$

2.4 Delta Encoding

We also looked into using delta encoding as part of our compression pipeline. The principle behind delta encoding is that instead of storing time sequenced data directly, the delta between subsequent time steps is calculated and saved. While the calculation of the delta itself will not lead to a decrease in data, it might reduce entropy, which can beneficially impact the performance of compression algorithms. We will formulate this more rigorously in the rest of this section.

We will start with a naive implementation of delta encoding and show that it is not suitable. We will use the following notation: At time step i , the full grid is denoted as g_i , and the

delta between the grids at time steps i and $i - 1$ is denoted as Δ_i . It is calculated by subtracting the previous grid from the current one: $\Delta_i = g_i - g_{i-1}$. Then, instead of passing the full grid sequence (g_1, g_2, \dots, g_N) to the dense data compressors, we use the delta sequence $(g_1, \Delta_2, \Delta_3, \dots, \Delta_N)$. The problem with this simple delta encoding approach is that the error starts to accumulate over time. To see this, we define

$$g'_i = g_i + \epsilon_i \tag{7}$$

$$\Delta'_i = \Delta_i + \eta_i = g_i - g_{i-1} + \eta_i \tag{8}$$

as the full frames and the delta frames which have gone through compression and therefore have an error ϵ_i and η_i applied to them. The grid at time step i that gets reconstructed by this is:

$$g_i^r = g'_1 + \sum_{j=2}^i \Delta'_j = g_i + \epsilon_1 + \sum_{j=2}^i \eta_j \tag{9}$$

Therefore the delta frame compression errors η_j start to accumulate over time. We can solve this problem rather easily. Instead of defining a delta frame from the “perfect” grid g_{i-1} , we use the reconstructed g_{i-1}^{r*} grid:

$$\begin{aligned} \Delta_i^{r*} &= \Delta_i^* + \eta_i = g_i - g_{i-1}^{r*} + \eta_i, \\ g_{i-1}^{r*} &= g'_1 + \sum_{j=2}^{i-1} \Delta_j^{r*} \end{aligned} \tag{10}$$

The g_i^{r*} can then be reformulated to

$$g_i^{r*} = g'_1 + \sum_{j=2}^i \Delta_j^{r*} = g_i + \eta_i \tag{11}$$

As we can see here the error does not start to accumulate over time, therefore this method is to be preferred over the naive one.

2.5 Implementation details

For the implementation, we had to take certain shortcuts and make adjustments to the theory in order to remain within reasonable computational limits. The primary simplification we made was to limit the possible number of grid points a trajectory may map to, for both the bipartite and the LP mapping, to a set of its closest neighbors. For bipartite mapping, this was the closest 200 neighbors, while for the LP mapping, it was the closest 50. However, this introduces the additional constraint of needing the starting locations of the trajectories to be evenly distributed. If they are not evenly distributed, i.e., if there are regions of low and high trajectory density, the mapping phase will fail, as we reach a point where we have locally more trajectories than possible grid points to which they may map. How limiting this constraint is depends on the type of research. For some studies, the initial configuration

is uniform, e.g., a global uniformity in Stoffels et al. (2025) Sprenger et al. (2017), and Bakels et al. (2025) or local uniformity in Pérez-Muñuzuri et al. (2018), and Wendisch et al. (2024). But there is also research such as Keune et al. (2022), Schielicke and Pfahl (2022), and Dey et al. (2023) where non-uniform initial configurations are used and which therefore cannot be handled by psit. A second shortcut we are taking is that we need more grid points than trajectories in practice, this is due to both considering only the closest neighbors and because the algorithmic implementation we chose to solve the minimal weight full bipartite matching problem performs much faster when there are more grid points compared to trajectories. In practice, we have around 50% more grid points than trajectories. These are the main considerations we had to take during implementation in order to strike a balance between computational performance and compression capabilities. In the future, it might be advisable to revisit these shortcuts, as removing them could lead to a broader applicable field of psit.

Some other minor implementation details are that we only implemented color encoding in combination with the JPEG 2000 compression algorithm, and that when compressing local rather than global trajectories (i.e., starting positions are bounded inside a box instead of spanning the entire globe), a longitude-latitude projection is used to map the grid points into 3D Cartesian space. This is because there is no simple inverse function for the presented projection method. Both of these differences do not have a large impact on the compression performance of psit, but could still be explored in future research.

3 Results

Previous research (Baker et al., 2014, 2016; Poppick et al., 2020) has already shown that simple error metrics are not enough to argue about the effectiveness of lossy compression algorithms in the case of weather and climate data. Therefore in this section we explore a set of experiments which should cover a range of different metrics and should give a general overview of how psit performs. In total we will carry out five different experiments:

1. Error metrics and comparison to ZFP (pyzfp v0.5.5) and SZ3 (v3.2.1) in order to compare performance to existing alternatives,
2. error value distribution in order to discern the creation of bias in the error,
3. trajectory density comparison against trajectories produced from perturbed wind fields in order to compare the impact of compression with the impact of data assimilation inaccuracies,

Table 1. The minimal and maximal values for the different data variables of the `tra_20200101_00` file.

Data variable	Minimum	Maximum
Pressure	32 hPa	974 hPa
Potential temperature	256 K	559 K
Potential vorticity	−50 PVU	66 PVU

Table 2. The minimal and maximal values for the different data variables of the `tra_20200101_00` file.

Data variable	Minimum	Maximum
Pressure	−1 hPa	1500 hPa
Temperature	182 K	313 K
Potential vorticity	−0.027 km ² kg ^{−1} s ^{−1}	0.039 km ² kg ^{−1} s ^{−1}
Specific humidity	−1.9 kg kg ^{−1}	0.022 kg kg ^{−1}

4. two case studies, one about warm conveyor belts and one about the Fukushima disaster, in order to see how psit performs in a more realistic environment,
5. and throughput and memory usage to observe the computational performance.

This should provide the reader with an insight into how psit behaves in different situations.

3.1 Input Data

During the following experiments we will use the following trajectory files:

- `tra_20200101_00`. A file produced using Lagranto (Wernli and Davies, 1997a; Sprenger and Wernli, 2015), based on ERA5, starting on 1 January 2020, 00:00 UTC and going until 2 January 2020, 00:00 UTC, with a total of 13 time steps (each two hours long). It consists of 5289 105 trajectories, distributed across 26 pressure levels at the first time step with horizontal spacing of 50 km, and vertically ranging from 50 to 550 hPa in 20 hPa increments. The additional data variables (ranges given in Table 1) it stores are potential temperature (TH) and potential vorticity (PV). The file is 1.3 GB in size.
- `tra_20200101_00_permuted`. This is the same trajectory file as `tra_20200101_00`, but with the trajectory order randomly permuted. Everything else remains the same. The order of the trajectories in the `tra_20200101_00` dataset follows the grid and pressure levels, which leads to continuity (i.e., neighbouring trajectories are similar) if the trajectories are traversed in order. By permuting them, this continuity is removed.
- `tra_20000101_00`. A trajectory file produced using Lagranto, based on ERA5, starting on 1 January 2000,

Table 3. Ablation study for psit on the `tra_20200101_00` trajectory file. For different configurations the file is compressed to a compression ratio of around 20 after which the RMSE error for the different data variables is considered. If we then pick the lowest value for each error column we can figure out the optimal configuration. It consists of using JPEG 2000 with delta encoding for all data variables except pressure where SZ3 with no delta encoding should be used, additionally we should use XYZ color encoding. The lowest value is always highlighted in a bold font.

Configuration	RMSE			
	lon–lat	<i>p</i>	TH	PV
JPEG 2000, no color, no delta	0.283	3.45	0.270	0.0942
SZ3, no color, no delta	0.613	1.59	0.256	0.116
JPEG 2000, no color, delta	0.161	2.12	0.170	0.0768
SZ3, no color, delta	2.38	1.95	0.197	0.162
JPEG 2000, XYZ color, delta	0.118	–	–	–

00:00 and extending 168 h until 8 January 2000 01:00. It consists of 5 867 016 trajectories distributed over the default 37 ERA5 pressure levels from 1000 to 1 hPa, with spacing of 40 km. The additional data variables traced along the trajectories are (ranges given in Table 2) temperature (*T*), potential vorticity (PV), and specific humidity (*Q*). The file is 23 GB in size.

- `tra_20000101_00_permuted`. The same trajectory file as `tra_20000101_00`, but with the trajectory order randomly permuted. The reasoning behind this is the same as for the `tra_20200101_00_permuted` file.

3.2 Performance of psit

There are different ways in which psit can be configured, leading to different compression behaviors. Here, we present one configuration stack that led to good results with the presented input data files. We found this optimal configuration by performing an ablation study in which we run psit in different configurations and then combine the best performing parts together. A summary of this can be seen in Table 3, where we compress the `tra_20200101_00` file in different configurations, but always with a ratio of around 20. We then look at the RMSE error of the different data variables in order to figure out what the best configuration is. In this configuration we are using the JPEG 2000 compression algorithm combined with delta encoding for every data variable except for pressure, where we use SZ3 (v3.2.1) with no delta encoding. Additionally, we utilize the XYZ color encoding scheme. This configuration is given by the configuration file in Listing 1, for a description of how these configuration files work, please refer to the user manual at (Pietak, 2025a). The compression factor, a metric that determines how strong the compression should be, is varied between the different runs and is the same for every data variable except for pressure,

```

1:  method:
2:      lon: jpeg
3:      lat: jpeg
4:      p: sz3
5:      TH: jpeg
6:      PV: jpeg
7:  delta_method:
8:      lon: r1
9:      lat: r1
10:     p: none
11:     TH: r1
12:     PV: r1
13:  color_method: xyz
14:  ratio:
15:     lon: <ratio>
16:     lat: <ratio>
17:     p: <1.5*ratio>
18:     TH: <ratio>
19:     PV: <ratio>
20:  bin: 0
21:  color_bits: 16
22:  num_workers: 26
23:  factor: 1.5
24:  exclude: BASEDATE

```

Listing 1. Configuration file for psit used for the runs presented in Sect. 3.2. The compression ratios are replaced by placeholders `<ratio>` to indicate that they are changed between the different runs, the compression ratio for the pressure variable is 1.5 larger than for the other ones.

where it is 1.5 times larger. We will use this configuration for all of the following experiments, except for the case study on the Fukushima disaster, as we are working with a different input file there.

3.2.1 Error metrics and comparison with ZFP and SZ3

In this experiment, a comparison between compression ratio and compression error for psit, ZFP, and SZ3 is conducted. The compression pipeline for ZFP works by treating each data variable as a 2D array over the time steps and the trajectories and compressing each of them individually, for ZFP we are using pyzfp v0.5.5 in fixed-accuracy mode. The pipeline for SZ3 is identical to the one used for ZFP, we are using the python interface pysz, version 3.2.1 and are running in the relative error bound mode. The error norms we use in this study are the L1, RMSE, L-infinity errors and the peak signal to noise ratio (PSNR) between the original and the compressed/decompressed data, which are obtained by calculating the errors defined by Eqs. (12), (13), (14), and (15) between each data variable of the original and compressed/decompressed data file.

$$l_1 = \frac{1}{n} \sum_{i=1}^n |x_i - y_i| \quad (12)$$

$$l_{\text{rmse}} = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - y_i)^2} \quad (13)$$

$$l_{\infty} = \max_{i=1, \dots, n} \{|x_i - y_i|\} \quad (14)$$

$$l_{\text{psnr}} = 20 \log_{10}(d) + 20 \log_{10}(l_{\text{rmse}}) \quad (15)$$

where n is the number of data points, x the original data, y the compressed/decompressed data and d the extent of the data. As we are using color encoding, we need to handle the longitude and latitude variables differently. Instead of using the difference, we use the central angle between the input and the output (defined as the angle on a great circle between two points on a sphere). The amount of compression error and therefore the compression ratio that is acceptable depends on the case in which the data is used. We therefore provide a wide range of compression ratios such that the reader is able to evaluate the general behavior and then based on this figure out what compression ratios are well suited for their needs. Such a comparison between psit and the other methods provides us with an overview of how psit compares to already

existing alternatives and therefore allows us to assess its viability.

In the first run, the `tra_20200101_00` and `tra_20200101_00_permuted` files are considered. An excerpt of this can be seen in Fig. 7, with the full plots given in Fig. B1 and a normalized version in Fig. B2 and the exact values in Tables B1, B2, B3, B4, and B5. When looking at both ZFP and SZ3, we see that the performance between the different input data files differs significantly. The performance of the permuted trajectory file is much worse compared to the non permuted one. This is because, for the `tra_20200101_00` file, the trajectories have been initialized on an equidistant grid and are stored in the file in an ordered manner. Therefore, by converting them into the 2D grids over time steps and trajectories, a naive smoothness (originating from spatial continuity) of the grid is created. When we permute the trajectories, this naive smoothness is destroyed, and the compression performance drops. This means that the compression performance of ZFP and SZ3 does not only depend on the data, but also on the way in which it is stored, something from which psit does not suffer. This also makes it difficult to predict and quantify the behavior of ZFP and SZ3 on input files. Comparing them with psit we see that psit gives consistent results independent on the input file. In the smaller compression range psit does not perform as well, this could be because during the mapping phase artificially more points are added. We also see that in general psit performs worse for pressure than it does for lon–lat this data variable dependent performance is something we quite often see.

In the second run, we consider the `tra_20000101_00` and `tra_20000101_00_permuted` files. This can be seen in Fig. B3 normalized version in Fig. B4, with exact values given in Tables B6, B7, B8, B9, and B10. For psit against ZFP we observe a similar, though less pronounced, behavior. Now, the difference between the two compression methods is less distinct, and psit only starts to perform better than ZFP for larger compression ratios. The performance of SZ3 now looks better and it outperforms psit for some cases. This decrease in performance for this larger trajectory file indicates that psit, in general, struggles with longer time ranges. A reason for this could be that over time, trajectories start to diverge, and therefore, the smoothness in the images, which is based on their initial position, gets lost (see e.g., Fig. 3).

From the above observations, we can make some general statements about the performance of psit. For smaller time ranges, psit performs on par with or better than ZFP and SZ3, with a notable exception being the pressure variable, due to its inherent noisy behavior. For longer time ranges, the performance of psit starts to degrade. Additionally, for cases where keeping a small L-infinity error is important, we should use SZ3 over JPEG 2000, as it bounds the L-infinity error. Moreover, the performance of ZFP and SZ3 is highly dependent on the way the input data is structured, something

from which psit does not suffer. We therefore conclude that psit is a viable alternative to ZFP and SZ3, and can deliver consistent results.

3.2.2 Error Value Distribution

The next experiment we are conducting is about the error distribution. For this, we compress the `tra_20200101_00_permuted` and `tra_20000101_00_permuted` files with psit, ZFP, and SZ3 and compression ratios of around 2.5, and 15, then calculate the difference from the original file and put normalized versions of these differences into a histogram. We chose these compression ratios as they represent both a small and a somewhat medium compression ratio. Based on the shape of the corresponding histogram we then can gauge if bias is introduced by the compression.

The distribution for the `tra_20200101_00_permuted` file is shown in Fig. 8, while the one for the `tra_20000101_00_permuted` is given in Fig. B5. For all variables we have that for psit the mean is nearly perfectly at 0. We also have that the shape for psit tends to be very similar over the different compression ratios. The other compression methods mostly also manage to keep the mean at zero, but tend to have vastly different (often non-smooth) shapes especially for the higher compression ratios. This indicates that psit generally does not introduce a bias.

3.3 Trajectory Density and Comparison to Perturbed Wind Fields

In this experiment, we examine the impact compression has on the position of the trajectories over time. To examine this impact, we calculate the densities of the trajectories starting at different pressure levels (1, 500, 1000 hPa) for increasing time steps (72, 120, 168 h). In addition to comparing uncompressed and compressed trajectories, we also look at uncompressed trajectories calculated from perturbed wind fields, simulating the impact of measurement uncertainties. In total, we consider the following trajectory datasets:

- Uncompressed trajectories,
- the `tra_20000101_00` trajectory file compressed with psit for compression ratios of 1.67, 12.8, 44.5,
- uncompressed trajectories (similar to `tra_20000101_00`) calculated from wind fields to which uniform noise to the horizontal/vertical components of magnitudes $0.01 \text{ m s}^{-1}/6.67 \times 10^{-4} \text{ Pa s}^{-1}$, $0.05 \text{ m s}^{-1}/3.33 \times 10^{-3} \text{ Pa s}^{-1}$, and $0.1 \text{ m s}^{-1}/6.67 \times 10^{-3} \text{ Pa s}^{-1}$ has been added.

The compression ratios are chosen in a way such that the entire range of compression ratios explored in Sect. 3.2.1 is covered. We then calculate 2D densities over longitude and

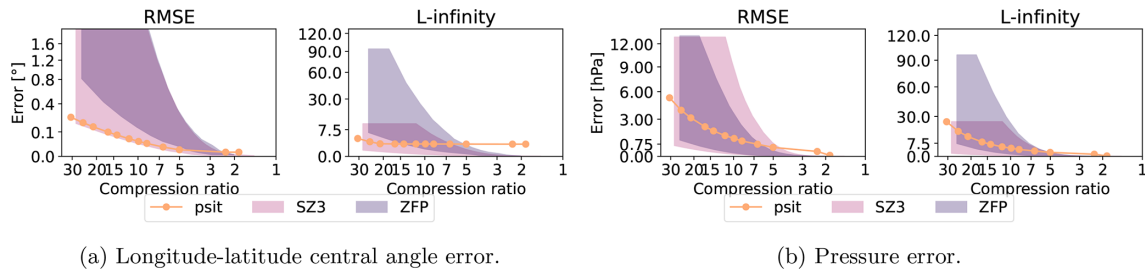


Figure 7. Excerpt of the comparison between psit, ZFP, and SZ3 for the `tra_20200101_00` and `tra_20200101_00_permuted` files. The configuration of our program is described in Sect. 3.2 and the corresponding config file is given in Listing 1. The RMSE and L-infinity error is compared to the achieved compression ratio for the central angle error (lon–lat) and pressure (p). In the plots the shaded area for ZFP and SZ3 corresponds to the range in compression performance between the two files. Note that psit performs the same for both, therefore, only one line is plotted.

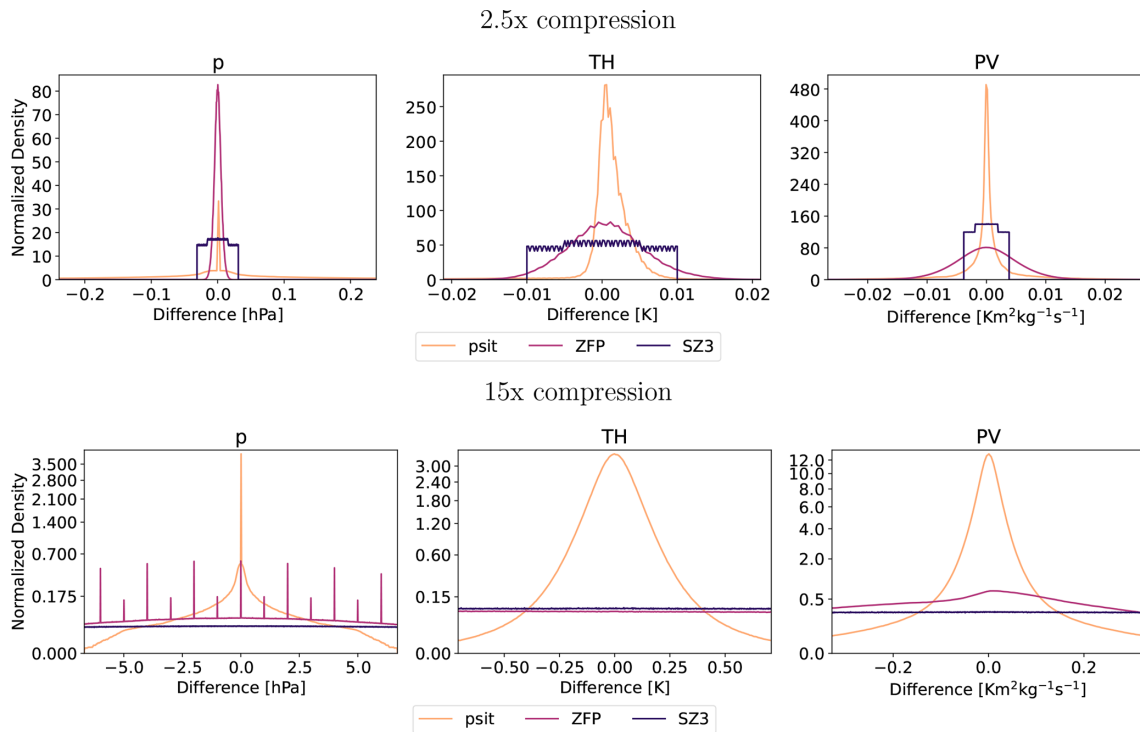


Figure 8. Area-normalized error distribution based on values for the `tra_20200101_00_permuted` dataset using the compression setup for psit of Sect. 3.2. We compress with compression ratios of 2.5, and 15. The area of the densities have been normalized and the x axis is cropped to 5 standard deviations of the psit distribution for all cases, except pressure with 2.5 times compression where it is 1 standard deviation. The data variables are pressure (p), potential temperature (TH), and potential vorticity (PV).

latitude (number trajectories over given area), and 1D densities over pressure (number trajectories over given pressure range), for which we consider the Jensen–Shannon divergence. We calculate the Jensen–Shannon divergence for the 2D case by partitioning the globe into equally sized bins and counting the number of trajectories per bin, and for the 1D case by using a histogram. This can be seen for example in Fig. 9. Such a density comparison between these different datasets and the uncompressed one allows us to see the im-

act compression has on the trajectory positions and also how compression behaves compared to perturbed wind fields.

The Jensen–Shannon divergence is shown in Tables B11, B12, and B13. From the tables, we can see that over time, the errors start to increase. The same (unsurprisingly) also holds for increasing compression ratio and perturbation magnitude. Furthermore, the error in pressure for the 500 and 1000 hPa pressure levels tends to be large and very similar between all runs, indicating that it is a very sensitive metric, as any perturbation leads to large errors. For the 2D diver-

gence metric, we find that with psit compressed trajectories in the compression range of 12.8 to 44.5 behave very similarly to perturbed wind fields with perturbation magnitudes of $0.05 \text{ m s}^{-1}/3.33 \times 10^{-3}$ to $0.1 \text{ m s}^{-1}/6.67 \times 10^{-3} \text{ Pa s}^{-1}$. The lowest compression ratio of 1.67 tends to outperform all the perturbed wind fields. This experiment demonstrates that psit behaves very similarly to perturbed wind fields, showing that the error incurred by compression lies within inaccuracies generated by data assimilation techniques.

Instead of looking at global trajectory densities, we can also look at local ones. For this, we take all the trajectories starting in the box from 0 to 15° E and 45 to 60° N , and plot their density alongside the difference to the uncompressed trajectories after 168 h (Fig. 10). From these plots, we can see that the general features are preserved under compression and that compression behaves very similarly to adding random perturbations to the wind field, again confirming that psit and perturbed wind fields lead to similar behavior of the trajectories.

3.4 Case Studies

3.4.1 Warm Conveyor Belts

The Warm Conveyor Belt (WCB) is one of three characteristic air streams in extratropical cyclones (Browning et al., 1973; Madonna et al., 2014). WCBs ascend in the cyclones' warm sector from the boundary layer to upper-tropospheric levels. This ascent, which occurs within 48 h and extends over 600 hPa, is associated with substantial cloud and precipitation formation (Pfahl et al., 2014). The WCB also influences the large-scale circulation and the weather predictability downstream of their interaction with the upper-level flow (Grams et al., 2011; Rodwell et al., 2018).

In order to carry out this experiment, we use the `tra_20000101_00_permuted` file and select all the trajectories which, in the first 48 h, rise by at least 600 hPa. We do this for uncompressed data and with psit compressed data (ratios of 1.67, 12.8, 44.5, exploring the compression range of 3.2.1). This is displayed in Fig. 11, with additionally a comparison with ZFP and SZ3 given in Fig. B6. For psit, we have that for all compression ratios, the general location and shape of the WCBs remain the same. With increasing compression ratio, the number of false positives increases faster than the number of false negatives, with these false positives mostly located in regions with already existing WCB activity. This asymmetric behavior is most likely due to the fact that the dense data compression algorithms perform smoothing of the data. Specifically, if a region with the majority of the selected trajectories gets smoothed out, the few unselected trajectories will start to behave like the selected ones, i.e., they will also be selected, leading to an increase in false positives. This observation is in line with the previous one that psit struggles with pressure while delivering good results for longitude and latitude. But still, in all cases, we see that

psit preserves the general structure, and only a very limited amount of false WCB activity gets generated due to compression (e.g., the region between the Southeast Atlantic Ocean and the Davis Sea).

3.4.2 The Fukushima accident

During the Fukushima Daiichi accident (IAEA, 2015) in 2011, radiation-contaminated material was released into the atmosphere. This leads to the question of where this radioactive material is transported and deposited (Wotawa, 2011; Chino et al., 2011; Katata et al., 2012; Yasunari et al., 2011; Stohl et al., 2012). In this experiment, we want to discuss the impact compression has on trajectories originating in the region of the Fukushima Daiichi power plant. We will not perform any detailed deposition analysis.

For this, we create 35'280 trajectories starting on 12 March 2011 at 00:00 in the box 140.74 to 141.3° E and 37.19 to 37.63° N over the pressure range of 764 to 996 hPa. The only information stored on the trajectories is their position, leading to a 59 MB file. For compression with psit, we choose JPEG 2000 with XYZ color and delta encoding. Note that here we compress trajectories that originate in a local region instead of a global one. We then compress the trajectories using psit, ZFP, and SZ3 with compression ratios around 6.5, 14, and 27. These compression ratios have been chosen to showcase both cases where compression works and cases where compression breaks down. We then plot the trajectories over time, as shown in Fig. 12. From the panels, we can see that for psit, ZFP, and SZ3 with increasing compression ratio, the results appear to get smoothed out and compression artifacts start to appear. We can also see that for psit the compression appears to perform worse, especially compared to Sect. 3.3. This is because the resulting grids have much smaller dimensions (48×37 grid points), leading to less data per grid, which results in poorer performance of the dense data compression algorithms. If we compare psit to the other two compression methods we can see that for lower compression ratios they are very similar to one another and only at the larger compression ratios can we observe psit delivering better results. This is in line with previous observations. This demonstrates that while psit is able to preserve the general structure and regions that would be affected, it is best used in the high ratio compression of files where many trajectories are present.

3.5 Throughput and Memory Usage

Until now, we have evaluated the performance of the program from an error standpoint. Here, we focus on performance from a computational perspective, discussing both throughput and memory usage. For this, the `tra_20000101_00` trajectory file is compressed in the same manner as presented in Sect. 3.2. This is performed on an AMD EPYC 7742 parallelized over 37 threads (for both compression and decom-

Trajectories starting at 500hPa

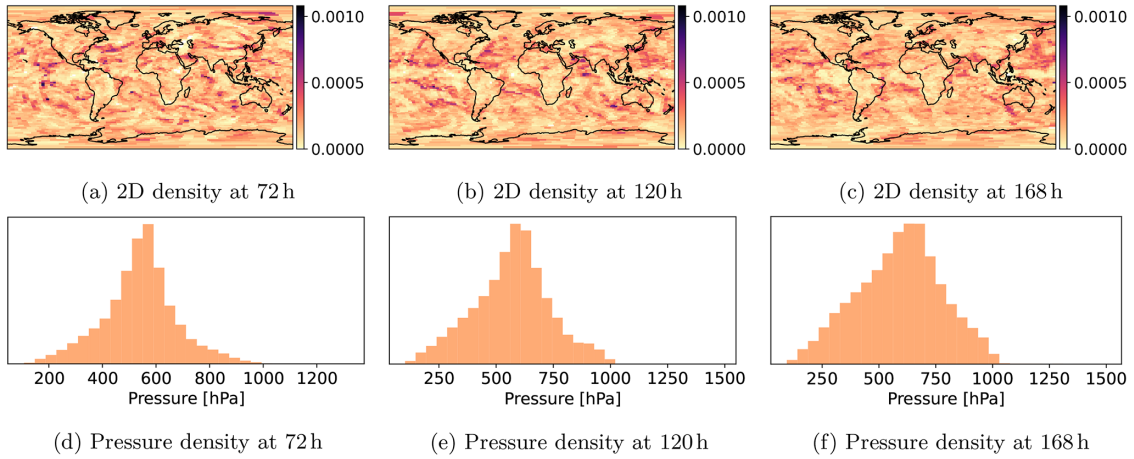


Figure 9. Reference image for how we calculate both the 1D and the 2D densities. For this we consider trajectories starting at 500 hPa and calculate the densities over multiple time steps. The 2D density is given over longitude and latitude, and as a 1D density over pressure.

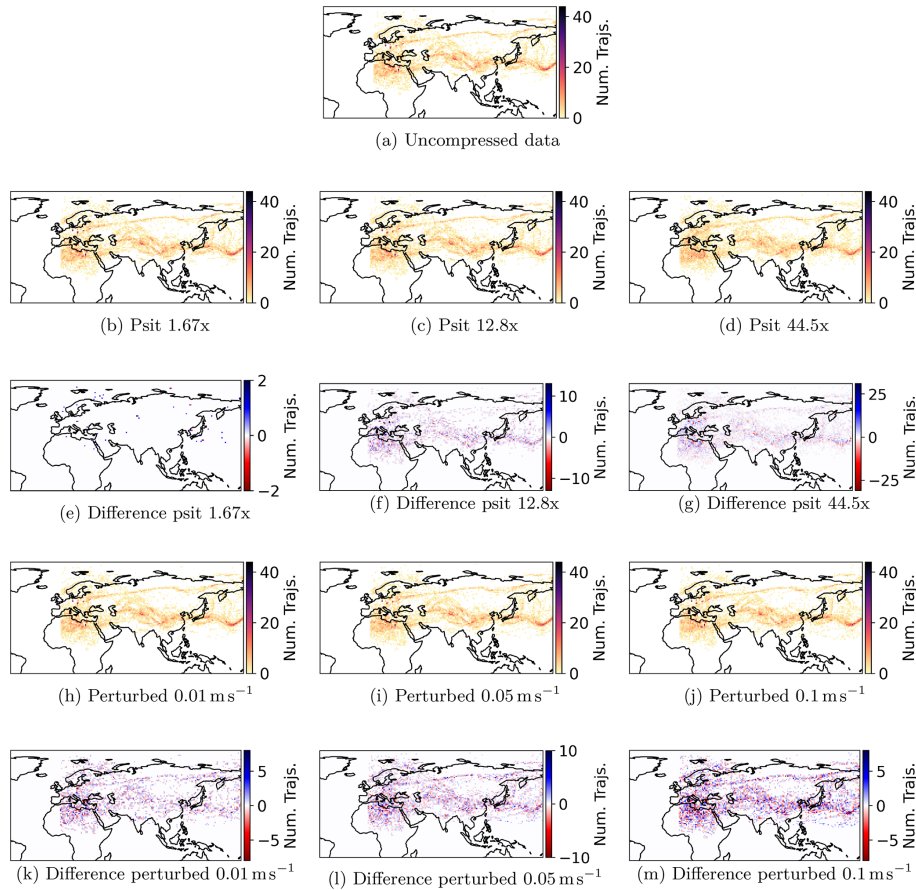


Figure 10. Comparison of the densities of trajectories starting in the box from 0 to 15° E times 45 to 60° N after 168 h. The trajectories we consider are uncompressed trajectories, compressed trajectories and trajectories calculated from perturbed wind fields.

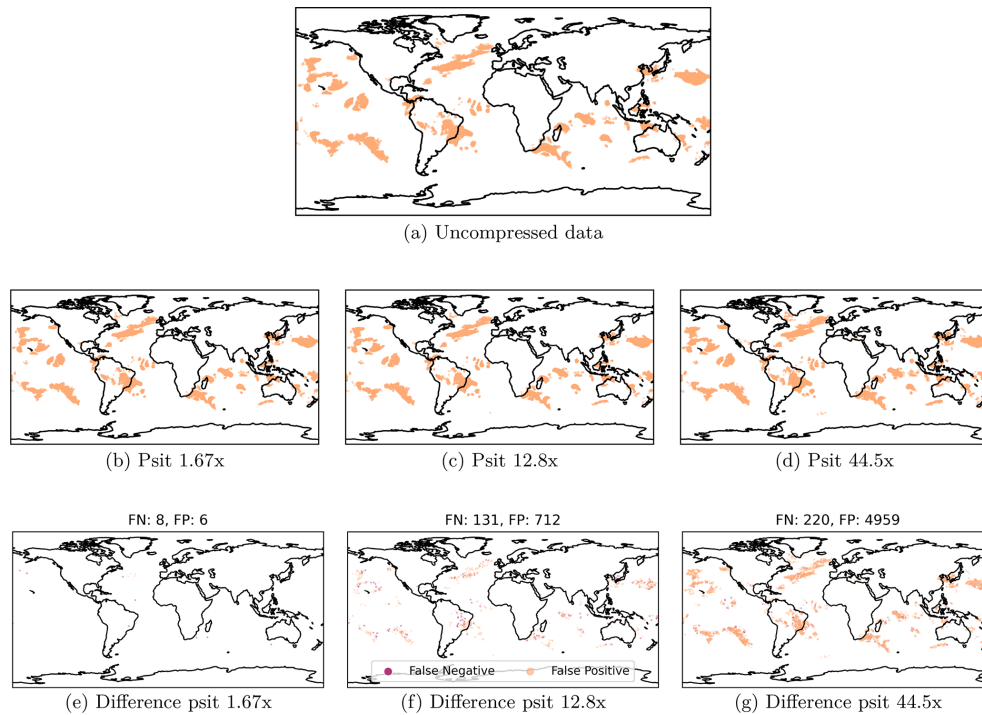


Figure 11. Warm Conveyor Belt calculations on the uncompressed data and the data compressed with psit at different compression ratios. Each point in a plot represents one trajectory which is part of a Warm Conveyor Belt (rises by at least 600 hPa in the first 48 h). Additionally the difference in term of false positive and false negatives is shown.

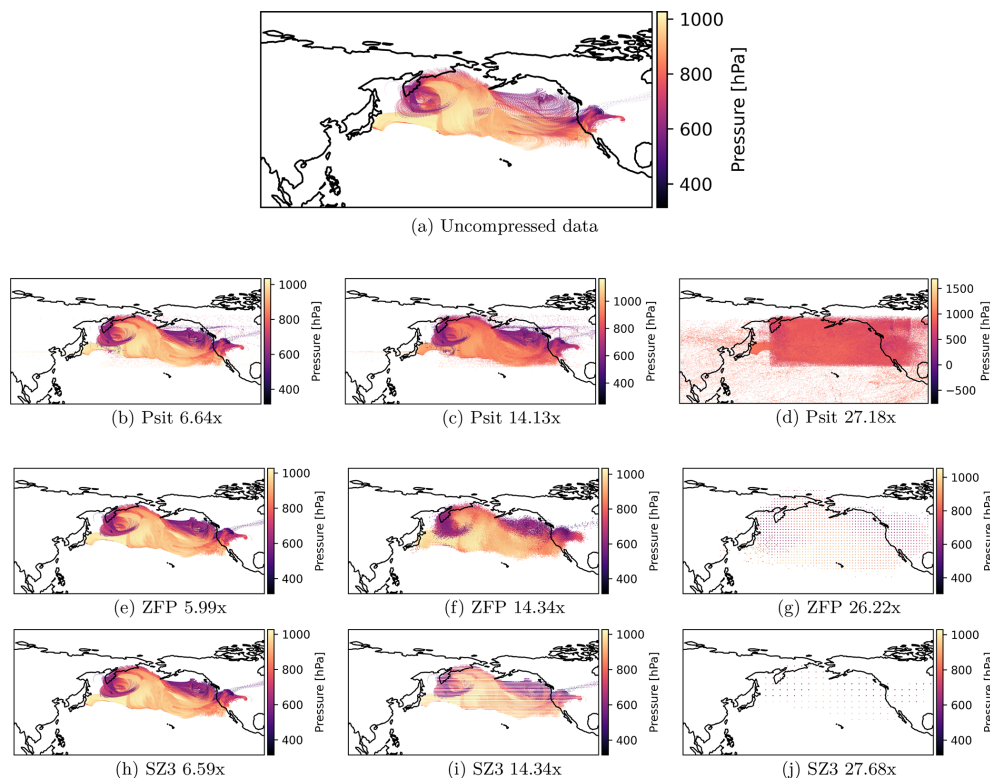


Figure 12. Comparison between trajectories starting at the site of the nuclear disaster at Fukushima for different compression ratios. On the top is the uncompressed data with, below it we have three rows of compressed data, one for psit, one for ZFP, and one for SZ3.

Table 4. Throughput for compression and decompression when running psit parallelized over 37 workers using a ramdisk.

Factor	Compression	Decompression
1	32.47 MB s ⁻¹	144.67 MB s ⁻¹
5	39.25 MB s ⁻¹	159.53 MB s ⁻¹
10	57.04 MB s ⁻¹	205.73 MB s ⁻¹
15	61.65 MB s ⁻¹	221.90 MB s ⁻¹
20	66.82 MB s ⁻¹	232.07 MB s ⁻¹
25	69.64 MB s ⁻¹	238.83 MB s ⁻¹
30	71.82 MB s ⁻¹	239.78 MB s ⁻¹
40	72.27 MB s ⁻¹	246.00 MB s ⁻¹
50	76.36 MB s ⁻¹	242.35 MB s ⁻¹
75	78.77 MB s ⁻¹	246.95 MB s ⁻¹
100	80.48 MB s ⁻¹	249.27 MB s ⁻¹
150	81.47 MB s ⁻¹	238.55 MB s ⁻¹

pression), using a ramdisk to filter out the impact of a slow filesystem. We then calculate the throughput in MB s⁻¹ for compression and decompression across the different compression ratios.

The results of the throughput experiments are presented in Table 4. The mean throughput of the compression stage is 65.67 MB s⁻¹, while the mean throughput of decompression is 222.14 MB s⁻¹. Decompression has a significantly larger throughput compared to compression because the expensive mapping stage only needs to be performed during compression. Another observation is that throughput increases as the compression factor gets larger. This is because a larger compression factor results in smaller file sizes, which, in turn, reduces the time needed to read and write data to disk. The memory usage for compression with factor 30 was 60GB. The part where most of the memory consumption comes from is in solving the minimal weight full bipartite matching, which is parallelized over multiple workers. Therefore, one can trade memory usage against runtime by decreasing the number of parallel workers.

4 Discussion

4.1 Applicable range and limitations

The main structural limitation of the program is the fact that the input data trajectories need to be uniformly distributed. This limitation originates from the way the mapping procedure has been implemented and has been touched upon in Sect. 2.5. The optimal fix for this limitation would be if the minimal weight full bipartite matching problem could be solved on a global level without having to restrict the grid points a trajectory may map to its closest 200 neighbors. A workaround for this is to artificially increase the ratio of grid points to trajectories, in this way, a mapping can be created at the cost of increased grid size, and because much of the data

in the grids is redundant, the compression algorithms might be able to compress these grids more effectively, leading to only a small increase in output size.

In addition to the structural limitations, there are also performance limitations, especially memory consumption and runtime. Here, the bottleneck is the solving of the minimal weight full bipartite matching problem in the mapping phase. The only way to solve this would be to either use a different mapping method or find a more efficient implementation of such an algorithm. Additionally, all the images are currently saved individually, leading to the creation of potentially thousands of files, which can be slow on certain filesystems. Furthermore, currently the compression pipeline is written in Python and was developed parallel to the exploration of methods, leading to sometimes non optimal implementations.

Based on the above discussed limitations of the implementation, we argue that the applicable range of psit is in the compression of uniformly distributed trajectories in an HPC environment.

4.2 Future Work

Of course, one obvious area of future work would be fixing the current limitations of psit. This would include removing the limitation of a trajectory only being able to map to a set of its closest neighboring grid points and generally improving the overall computational performance of the pipeline in terms of runtime and memory usage. Fixing these limitations would be quite beneficial for psit, as it would allow it to be used in a wider range of applications.

The mapping method currently used by psit is based on graph theory, specifically the solution to a bipartite matching problem. However, there are likely other methods by which this could be done. For instance, instead of solving the problem directly in one step, an initial simple mapping (such as via nearest neighbors) could be created, which is then iteratively refined to fulfill the initial constraints, with the goal being to create a valid mapping that is also smooth. Further exploration of what “smooth” means could also be valuable. Currently, smoothness is achieved by using the trajectory positions at the first time step. What would happen if the mean position were used instead? Alternatively, could the smoothness of a grid be quantized to create a mapping that aims to optimize it? As the mapping phase is integral to the entire pipeline and currently presents the performance bottleneck, exploring these questions would be greatly beneficial.

The methods we use to describe compression performance are error metrics, comparison with perturbed/bitrounded wind fields, and real world examples in the form of two case studies. However, there are some intricacies to consider. Error metrics provide a straightforward way to quantify the incurred error, but they often lack a strong connection to real-world applications. It can be challenging to predict the impact of compression on real-world research based solely on a

simple error metric. Is an RMSE error of 0.5 hPa in pressure acceptable, or should it be less than 0.3 hPa? Moreover, simple error metrics do not illustrate the relationship between different variables. How does the error in the longitude variable compare to the error in the latitude variable? Are they similar, or is there more error in one than the other? Does the error in the temperature variable depend on the pressure variable? A statistical analysis could examine correlations between various data variables, but this can quickly become tedious while leaving the fundamental question unanswered of how much correlation is acceptable. On the other hand, case studies simulate scenarios where the data might be used. We believe this approach establishes a better foundation for understanding the impact of compression. However, the challenge is that one can only analyze singular examples, and the generalizability of findings from one test to another remains an open question. While writing this paper, we spent considerable time discussing with experts in the field how such errors could be described and interpreted, there we found three major points that we think are important: compression should not change the outcome of experiments, compression errors should be smaller than measurement errors, compression should not influence the timings and occurrence of physical processes (e.g., cloud formation). But even these three points are not universally applicable, e.g., how does one compare measurement errors to compression errors, which physical processes are important enough to be conserved and to which degree should they be conserved. We tried to address these challenges in this paper, but we think that it would be beneficial and interesting to dedicate time into researching the possibilities of developing a benchmark that can help in understanding the impact of these errors. This benchmark could combine error metrics with a list of real-world applications, providing users with a general overview of how compression performs across different aspects.

5 Conclusions

The initial problem we faced and aimed to solve was that increasing computational power leads to an increase in data, making storage unfeasible. Therefore, a method needs to be devised to reduce it, with one approach being compression. While compression schemes exist for the Eulerian frame, no equivalent alternatives exist for the Lagrangian one. We therefore developed psit, a system to compress Lagrangian flows, to address this issue. In this paper, we presented how psit works and evaluated its performance using error metrics and case studies. In these, we demonstrated that in most cases, compression performance equivalent to or superior to ZFP and SZ3 can be achieved. We showed that the trajectory density of compressed trajectories (ratio 30 to 40) after 168 h behaves similarly to uncompressed trajectories calculated from perturbed wind fields (with a perturbation magnitude of $0.1 \text{ m s}^{-1}/6.67 \times 10^{-3} \text{ Pa s}^{-1}$). Additionally, we an-

alyzed the impact of compression on subsequent research through two case studies: warm conveyor belts and fallout prediction of the Fukushima accident. While psit imposes some limitations on the input data, namely requiring a uniform initial distribution, it can be used for general purpose compression of Lagrangian weather data and can serve as a foundation for further research in this area.

Appendix A: Proof of Total Unimodularity

In this section, we provide a proof that the LP presented in the LP mapping of Sect. 2.2 is an integral LP. An LP of the form $\{\min \mathbf{c}^T \mathbf{x} \mid \mathbf{A}\mathbf{x} \geq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ is integral if the constraint matrix \mathbf{A} is totally unimodular and the right hand side vector \mathbf{b} is integral. In this mapping, the \mathbf{b} vector consists solely of values equal to 1, making it integral. The only remaining step is to show that the matrix \mathbf{A} is totally unimodular.

Before we begin the proof, we must discuss the structure of the matrix \mathbf{A} . In the definition of the LP, it was stated that the vector \mathbf{x} consists of the values $x_{t,p}$ with $t \in T$ and $p \in P$, and that there are two types of constraints:

$$\sum_{t \in N^-(p)} x_{t,p} = 1 \quad \forall p \in P \quad (\text{A1})$$

$$\sum_{p \in N^+(t)} x_{t,p} \geq 1 \quad \forall t \in T \quad (\text{A2})$$

Note that in contrast to the definition given in Sect. 2.2, here we also have the terms $N^-(p)$ and $N^+(t)$ which denote the possible matchings of trajectories to grid points, i.e., $N^-(p)$ for all the trajectories which may map to a point p and $N^+(t)$ for all the grid points which can be mapped by a trajectory t . This is such that the proof is also valid for the shortcuts we took during implementation. To transform these constraints into canonical form, we need to rewrite the equality constraint as a set of two inequalities. The relationship $a = b$ can be expressed as $a \geq b$ and $-a \geq -b$. Thus, the constraints become:

$$\sum_{t \in N^-(p)} x_{t,p} \geq 1 \quad \forall p \in P \quad (\text{A3})$$

$$\sum_{t \in N^-(p)} -x_{t,p} \geq -1 \quad \forall p \in P \quad (\text{A4})$$

$$\sum_{p \in N^+(t)} x_{t,p} \geq 1 \quad \forall t \in T \quad (\text{A5})$$

Next, we need to express these constraints in matrix form. To do this, we must establish a representation for the vector \mathbf{x} . Without loss of generality, we can assume that \mathbf{x} is constructed by first looping over all trajectories and then looping over all grid points to which each trajectory may map. This means that blocks of $x_{t,p}$ with the same t and varying p are stacked on top of each other. For example, if we have three trajectories t_1, t_2, t_3 and four grid points p_1, p_2, p_3, p_4 , the

resulting vector \mathbf{x} could look like this:

$$\mathbf{x} = \begin{bmatrix} x_{I_1, p_2} \\ x_{I_1, p_3} \\ x_{I_2, p_1} \\ x_{I_2, p_3} \\ x_{I_3, p_4} \end{bmatrix} \tag{A6}$$

With the shape of the vector \mathbf{x} established, we can now construct the matrix \mathbf{A} . This matrix consists of two parts. In the first part, called \mathbf{A}_1 , we encode the first two constraints (Eqs. A3 and A4), and in the second part, called \mathbf{A}_2 , we encode constraint Eq. (A5).

For the matrix \mathbf{A}_1 , each row corresponds to a grid point. A row resulting from constraint Eq. (A3) consists of 1s and 0s, with at most one 1 per block of the same trajectory in \mathbf{x} . The exact distribution of these 1s depends on the shape of \mathbf{x} , but what is important is that no two rows have a 1 in the same column. This follows directly from the structure of \mathbf{x} and the fact that each row corresponds to a grid point. Another important property is that when constraints Eqs. (A3) and (A4) are combined, there are always pairs of rows that are identical, except that the sign of one row is the opposite of the other. Using the same example as above, the matrix \mathbf{A}_1 will look like this:

$$\mathbf{A}_1 = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & -1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -1 \end{bmatrix} \tag{A7}$$

Here, the pairs of rows with opposite signs are clearly visible.

Matrix \mathbf{A}_2 results from constraint Eq. (A5). In this case, each row corresponds to a trajectory, and all columns are set to 1 where a corresponding $x_{t,p}$ variable exists for that specific trajectory. Due to the construction of the \mathbf{x} vector, all the different trajectories are grouped into blocks, and for each block, all the grid points that the trajectory can map to are indicated. Consequently, each row will contain a continuous sequence of 1s corresponding to the $x_{t,p}$ values of the \mathbf{x} vector with a specific t . This matrix will have a staircase structure, ensuring that no two rows have a 1 in the same column. Combining the previous part of \mathbf{A} with this new one, it can be defined as follows:

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & -1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -1 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{A8}$$

The last step is to prove that this matrix is totally unimodular. The definition of total unimodularity states that for every square submatrix \mathbf{S} of \mathbf{A} , its determinant must be either -1 , 0 , or 1 . Proving this directly can be quite challenging. However, Ghouila-Houri has shown that a matrix is totally unimodular if and only if for every subset R of rows of the matrix, there exists an assignment $s : R \rightarrow \pm$ of signs such that, when summed, every element in the resulting vector is either -1 , 0 , or 1 . Therefore, we must find such a mapping $s : R \rightarrow \pm$ to conclude the proof.

To define this mapping, we use the property that in both parts of the matrix, no two rows have a 1 with the same sign in the same column. Initially, we assume that constraint Eq. (A4) does not exist, leading to the following definition of the mapping:

$$\forall r \in R : s(r) = \begin{cases} + & \text{if } r \in \mathbf{A}_1 \\ - & \text{if } r \in \mathbf{A}_2 \end{cases} \tag{A9}$$

We will first sum the vectors separately for the two parts, \mathbf{A}_1 and \mathbf{A}_2 . Due to the previously mentioned property, the vector derived from \mathbf{A}_1 will contain only elements from $\{0, 1\}$, while the vector from \mathbf{A}_2 will consist only of elements from $\{-1, 0\}$. The sum of these two vectors will thus produce a vector with elements from $\{-1, 0, 1\}$, confirming that a valid mapping has been found.

Next, we will reintroduce constraint Eq. (A4) into the formulation. The mapping can be adjusted to accommodate this case by making a small extension: a row $r \in R$ resulting from Eq. (A4) will be mapped to a $-$ if and only if the same row with the opposite sign is *not* in R . In this way, two scenarios can occur: either both the positive and negative variants of a row are in R , in which case they will cancel out, or only one of them, either the negative or the positive variant, is in R . In the latter case, a $-$ would be assigned to the negative variant, transforming the problem into the same form as if constraint Eq. (A4) did not exist at all. Therefore, it has been demonstrated that for any arbitrary subset R of rows from \mathbf{A} , a mapping $s : R \rightarrow \pm$ has been found that makes the matrix totally unimodular, thereby proving that the LP is integral.

Appendix B: Additional performance data of psit

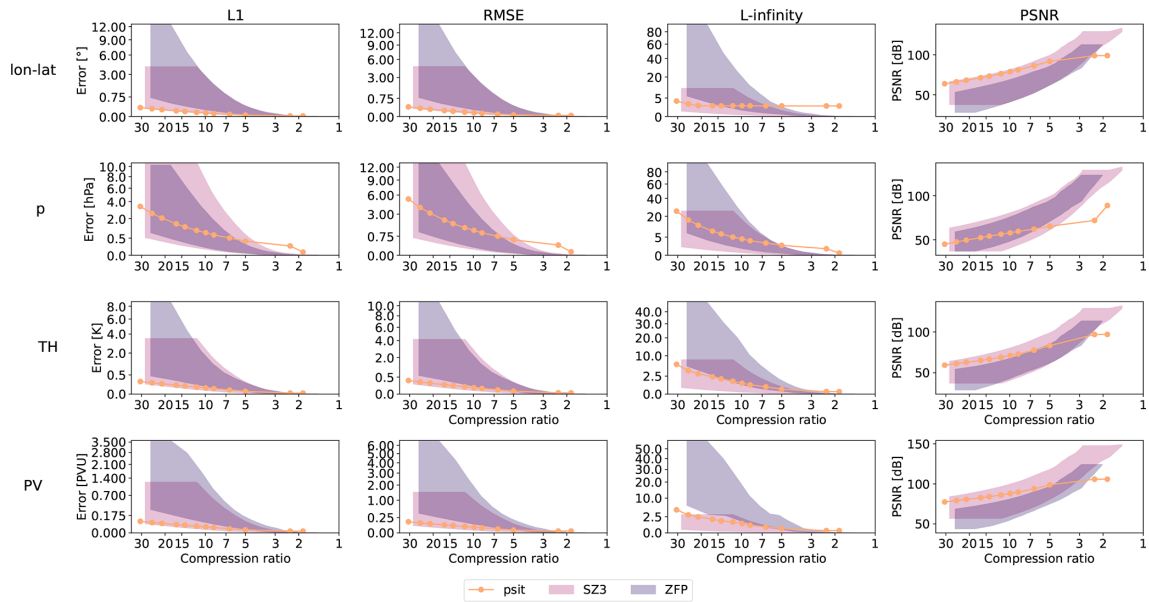


Figure B1. Comparison between psit, ZFP, and SZ3 for the `tra_20200101_00` and `tra_20200101_00_permuted` files. The configuration of our program is described in Sect. 3.2 and the corresponding config file is given in Listing 1. The L1, RMSE, L-infinity error and PSNR is compared to the achieved compression ratio for the central angle error (lon-lat) and the other data variables (pressure (p), potential temperature (TH), and potential vorticity (PV)). In the plots the shaded area for ZFP and SZ3 corresponds to the range in compression performance between the two files. Note that psit performs the same for both, therefore, only one line is plotted.

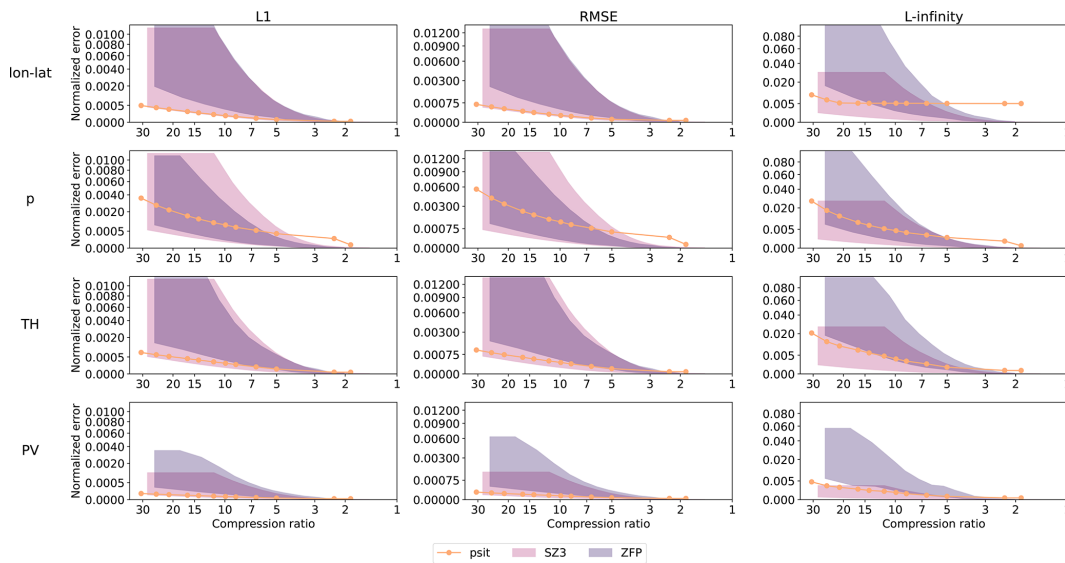


Figure B2. Data variable normalized error comparison between psit, ZFP, and SZ3 for the `tra_20200101_00` and `tra_20200101_00_permuted` files. The configuration of our program is described in Sect. 3.2 and the corresponding config file is given in Listing 1. The L1, RMSE, and L-infinity errors are normalized and compared to the achieved compression ratio for the central angle error (lon-lat) and the other data variables (pressure (p), potential temperature (TH), and potential vorticity (PV)). In the plots the shaded area for ZFP and SZ3 corresponds to the range in compression performance between the two files. Note that psit performs the same for both, therefore, only one line is plotted.

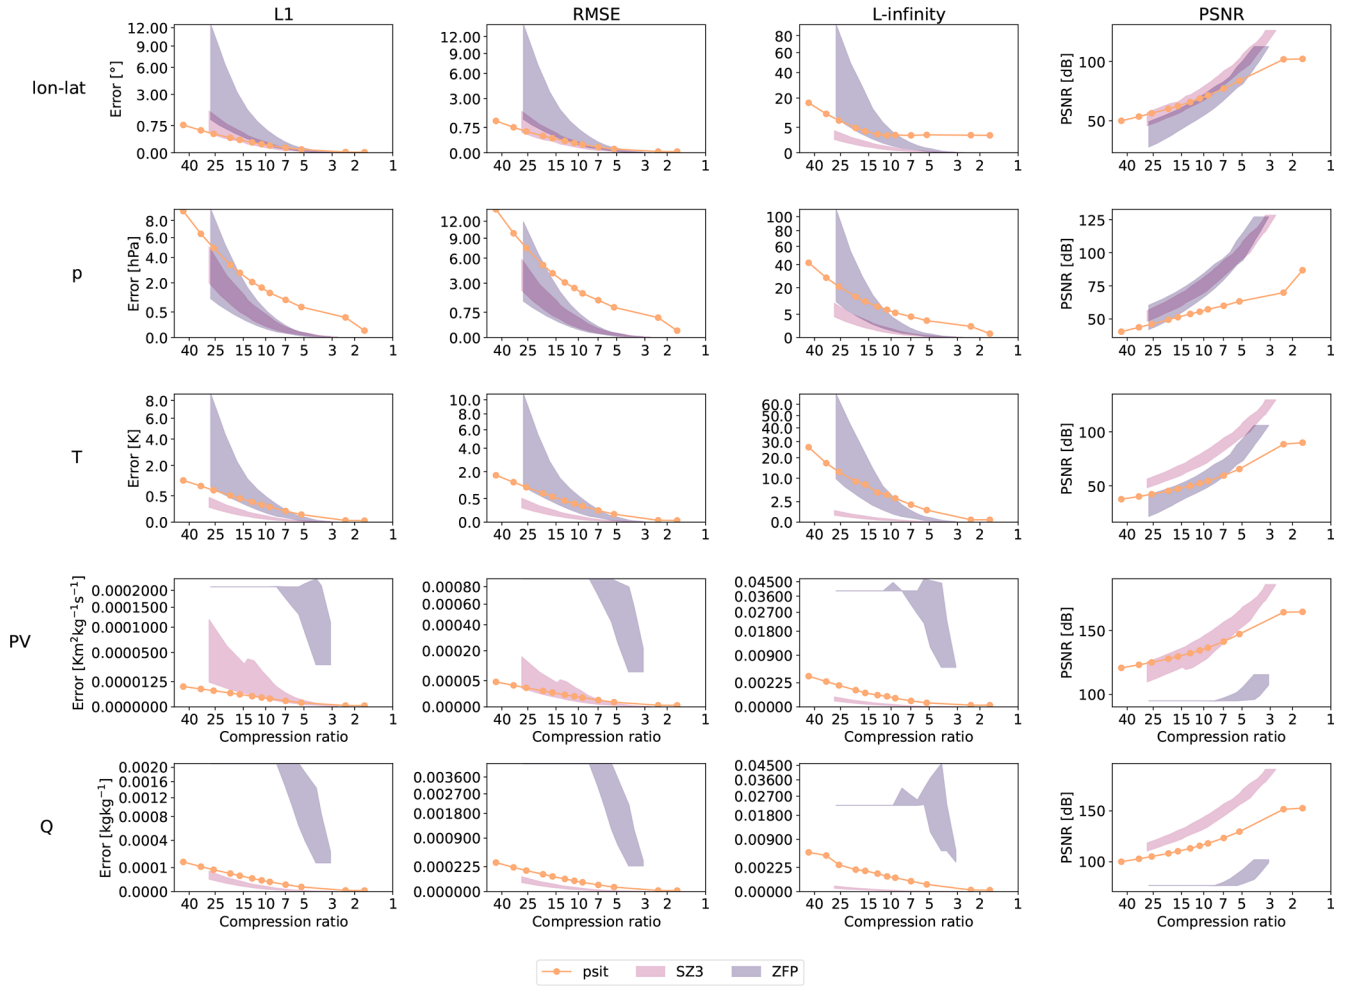


Figure B3. Comparison between psit, ZFP, and SZ3 for the `tra_20000101_00` and `tra_20000101_00_permuted` files. The configuration of our program is described in Sect. 3.2 and the corresponding config file is given in Listing 1. The L1, RMSE and L-infinity error and PSNR is compared to the achieved compression ratio for the central angle error (lon–lat) and the other data variables (pressure (p), temperature (T), potential vorticity (PV), and specific humidity (Q)). In the plots the shaded area for ZFP and SZ3 corresponds to the range in compression performance between the two files. Note that psit performs the same for both, therefore, only one line is plotted.

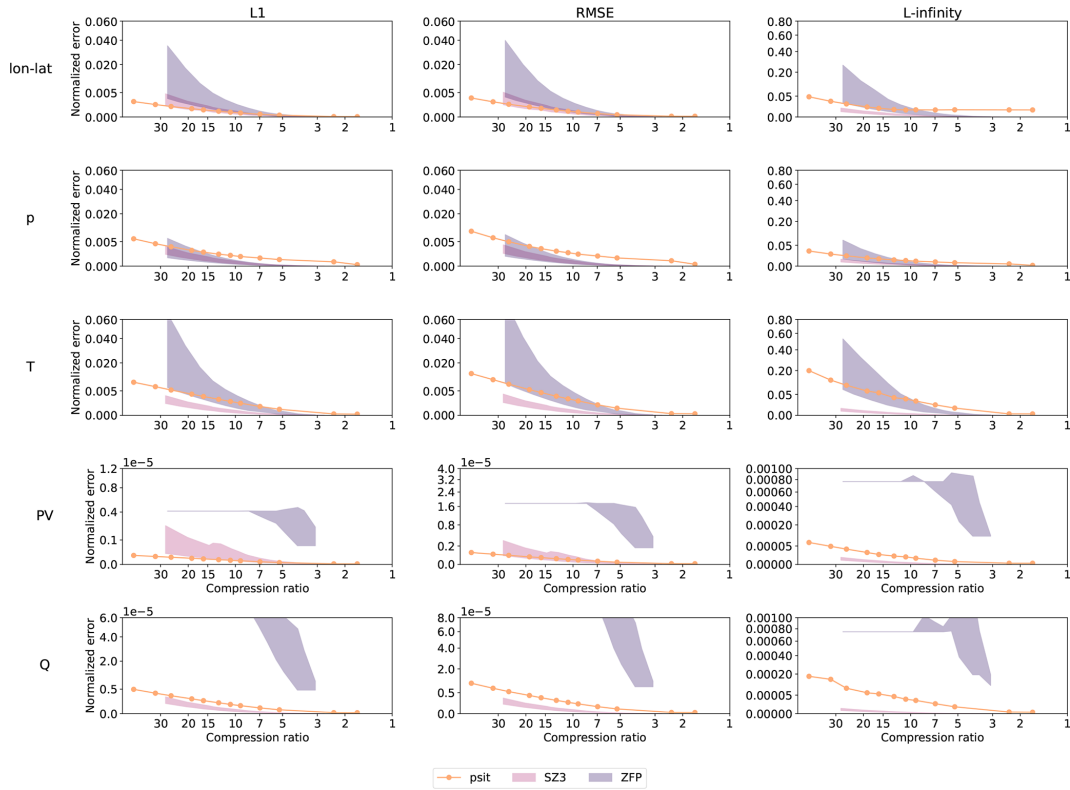


Figure B4. Data variable normalized error comparison between psit, ZFP, and SZ3 for the tra_20000101_00 and tra_20000101_00_permuted files. The configuration of our program is described in Sect. 3.2 and the corresponding config file is given in Listing 1. The L1, RMSE, and L-infinity errors are normalized and compared to the achieved compression ratio for the central angle error (lon–lat) and the other data variables (pressure (p), temperature (T), potential vorticity (PV), and specific humidity (Q)). In the plots the shaded area for ZFP and SZ3 corresponds to the range in compression performance between the two files. Note that psit performs the same for both, therefore, only one line is plotted.

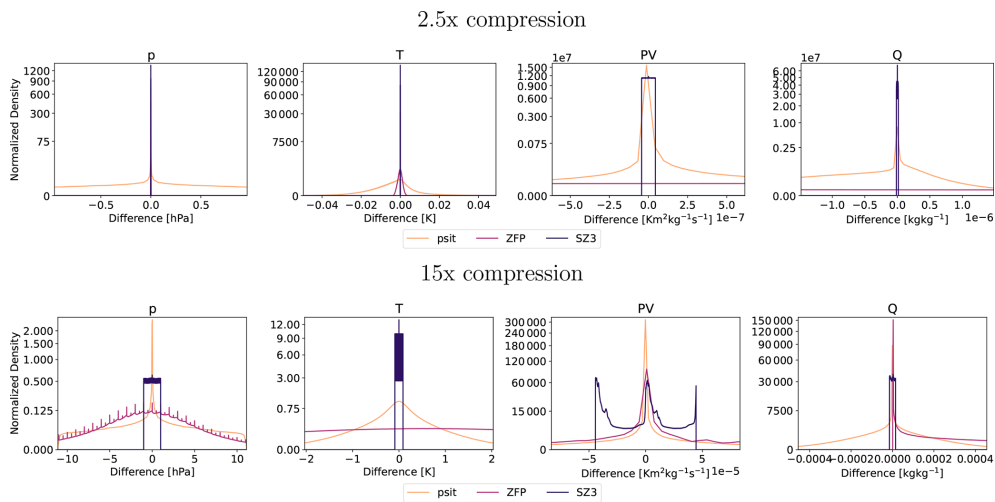


Figure B5. Area-normalized error distribution based on values for the tra_20000101_00_permuted dataset using the compression setup of Sect. 3.2 for psit with compression ratios of 2.5, and 15. The area of the densities have been normalized and the x axis is cropped to 2 standard deviations of the psit distribution for the 2.5 times compression case and 5 standard deviations for the 15 times compression case. The data variables are pressure (p), temperature (T), potential vorticity (PV), and specific humidity (Q).

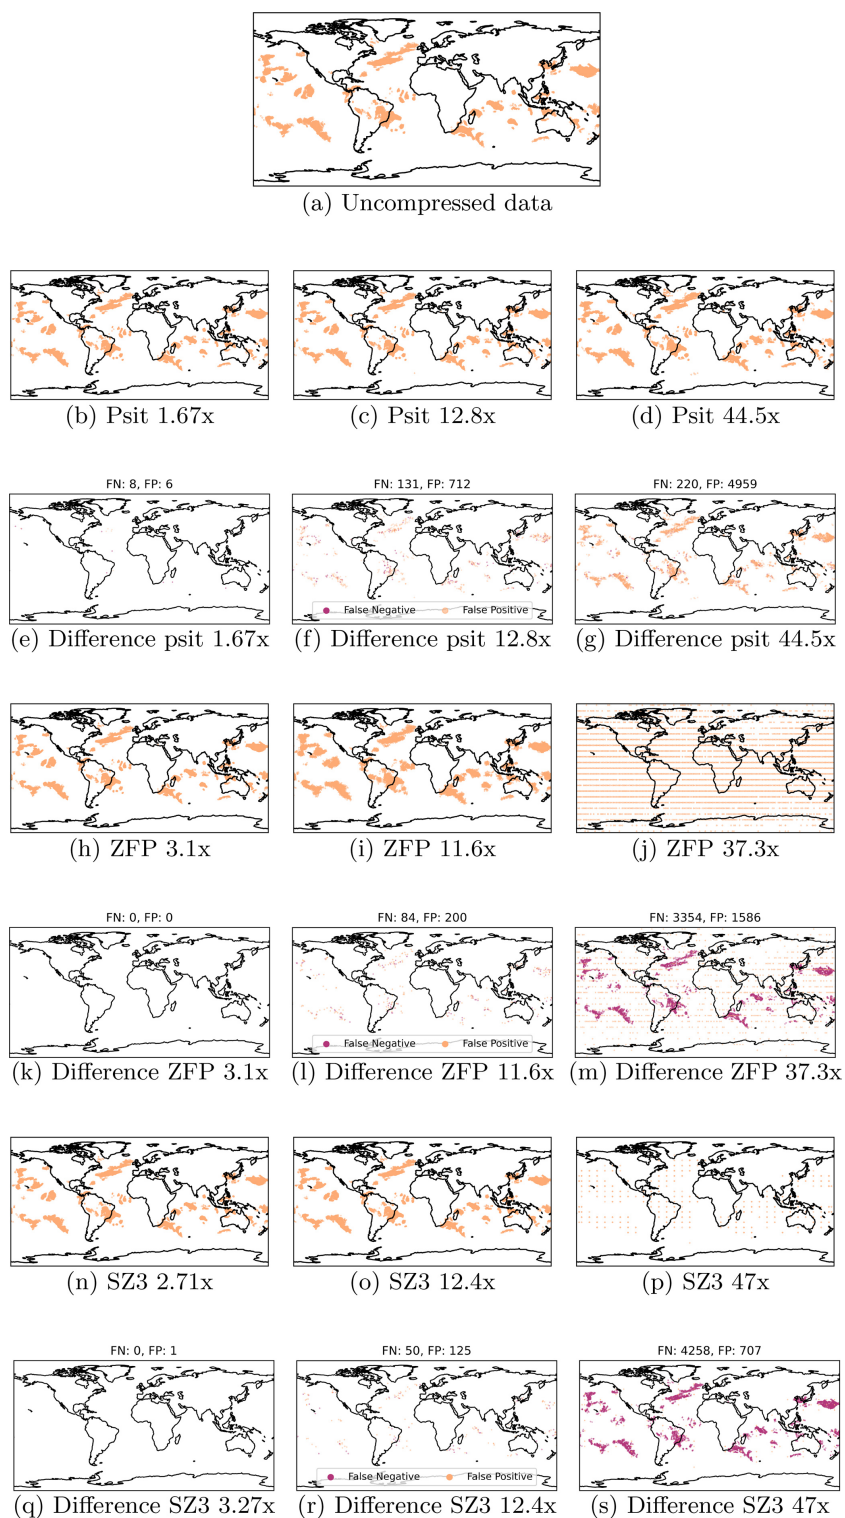


Figure B6. Warm Conveyor Belt calculations on the uncompressed data and the data compressed with different compression ratios and compressors. Each point in a plot represents one trajectory which is part of a Warm Conveyor Belt (rises by at least 600 hPa in the first 48 h). Additionally the difference in term of false positive and false negatives is shown.

Table B1. Error values and compression ratios for psit on the tra_20200101_00 dataset.

Factor	Ratio	lon-lat			<i>p</i>			TH			PV		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
1	1.86	1.06e-3	4.03e-3	1.80e+0	2.37e-2	3.41e-2	1.12e-1	2.21e-3	4.27e-3	6.84e-2	2.72e-3	5.33e-3	7.17e-2
5	2.32	1.08e-3	4.06e-3	1.80e+0	1.66e-1	2.39e-1	7.82e-1	2.37e-3	4.35e-3	7.08e-2	2.77e-3	5.35e-3	7.19e-2
10	5.01	6.48e-3	9.19e-3	1.82e+0	3.53e-1	5.09e-1	1.68e+0	1.58e-2	2.11e-2	2.35e-1	7.43e-3	1.21e-2	2.31e-1
15	6.58	1.33e-2	1.73e-2	1.84e+0	5.12e-1	7.40e-1	2.46e+0	3.04e-2	4.03e-2	4.85e-1	1.27e-2	2.13e-2	3.91e-1
20	8.63	2.50e-2	3.14e-2	1.84e+0	6.85e-1	9.95e-1	3.35e+0	5.17e-2	6.85e-2	7.57e-1	2.04e-2	3.44e-2	6.56e-1
25	9.91	3.28e-2	4.07e-2	1.86e+0	8.29e-1	1.21e+0	4.13e+0	6.46e-2	8.59e-2	1.01e+0	2.49e-2	4.21e-2	9.34e-1
30	11.61	4.44e-2	5.49e-2	1.86e+0	9.86e-1	1.45e+0	5.03e+0	8.16e-2	1.09e-1	1.39e+0	3.09e-2	5.20e-2	1.21e+0
40	14.20	6.31e-2	7.78e-2	1.85e+0	1.26e+0	1.87e+0	6.70e+0	1.05e-1	1.42e-1	1.87e+0	3.90e-2	6.57e-2	1.45e+0
50	16.48	7.93e-2	9.82e-2	1.86e+0	1.51e+0	2.26e+0	8.38e+0	1.25e-1	1.70e-1	2.35e+0	4.56e-2	7.68e-2	1.84e+0
75	21.10	1.12e-1	1.41e-1	1.90e+0	2.06e+0	3.14e+0	1.25e+1	1.61e-1	2.23e-1	3.13e+0	5.79e-2	9.75e-2	2.44e+0
100	25.01	1.38e-1	1.78e-1	2.50e+0	2.55e+0	3.93e+0	1.68e+1	1.90e-1	2.66e-1	4.08e+0	6.79e-2	1.15e-1	3.02e+0
150	30.62	1.77e-1	2.34e-1	3.60e+0	3.36e+0	5.30e+0	2.51e+1	2.32e-1	3.31e-1	6.27e+0	8.29e-2	1.41e-1	4.76e+0

Table B2. Error values and compression ratios of the ZFP baseline run on the tra_20200101_00 dataset. Note that for some different tolerances the compression ratios are the same, this is probably because of the grouping into “bit planes” that is done by the embedded coding strategy of ZFP.

Tolerance	Ratio	lon-lat			<i>p</i>			TH			PV		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
0.01	2.90	7.17e-4	8.39e-4	4.13e-3	5.11e-4	6.52e-4	3.51e-3	5.11e-4	6.52e-4	3.51e-3	5.11e-4	6.52e-4	3.58e-3
0.05	3.61	2.81e-3	3.27e-3	1.58e-2	2.05e-3	2.61e-3	1.44e-2	2.04e-3	2.61e-3	1.39e-2	2.03e-3	2.59e-3	2.23e-2
0.1	4.10	5.47e-3	6.36e-3	3.08e-2	4.09e-3	5.21e-3	2.86e-2	4.09e-3	5.21e-3	2.91e-2	4.01e-3	5.11e-3	4.13e-2
0.5	6.64	3.59e-2	4.16e-2	2.47e-1	3.25e-2	4.14e-2	2.38e-1	3.12e-2	3.98e-2	2.25e-1	2.58e-2	3.37e-2	3.66e-1
1.0	8.11	6.51e-2	7.57e-2	4.22e-1	6.45e-2	8.21e-2	4.67e-1	5.83e-2	7.44e-2	4.30e-1	4.41e-2	5.86e-2	7.40e-1
2.0	10.13	1.17e-1	1.36e-1	7.60e-1	1.26e-1	1.61e-1	9.06e-1	1.04e-1	1.33e-1	8.50e-1	7.24e-2	9.77e-2	1.44e+0
3.0	10.13	1.17e-1	1.36e-1	7.60e-1	1.26e-1	1.61e-1	9.06e-1	1.04e-1	1.33e-1	8.50e-1	7.24e-2	9.77e-2	1.44e+0
4.0	13.08	2.05e-1	2.40e-1	1.55e+0	2.44e-1	3.11e-1	1.77e+0	1.79e-1	2.30e-1	1.66e+0	1.16e-1	1.58e-1	2.85e+0
5.0	13.08	2.05e-1	2.40e-1	1.55e+0	2.44e-1	3.11e-1	1.77e+0	1.79e-1	2.30e-1	1.66e+0	1.16e-1	1.58e-1	2.85e+0
7.0	13.08	2.05e-1	2.40e-1	1.55e+0	2.44e-1	3.11e-1	1.77e+0	1.79e-1	2.30e-1	1.66e+0	1.16e-1	1.58e-1	2.85e+0
10.0	17.66	3.63e-1	4.23e-1	3.08e+0	4.58e-1	5.88e-1	3.52e+0	2.98e-1	3.84e-1	3.18e+0	1.88e-1	2.51e-1	3.21e+0
15.0	17.66	3.63e-1	4.23e-1	3.08e+0	4.58e-1	5.88e-1	3.52e+0	2.98e-1	3.84e-1	3.18e+0	1.88e-1	2.51e-1	3.21e+0
20.0	25.54	7.03e-1	8.14e-1	6.08e+0	8.31e-1	1.08e+0	7.29e+0	4.57e-1	5.92e-1	5.63e+0	2.98e-1	3.98e-1	6.59e+0
25.0	25.54	7.03e-1	8.14e-1	6.08e+0	8.31e-1	1.08e+0	7.29e+0	4.57e-1	5.92e-1	5.63e+0	2.98e-1	3.98e-1	6.59e+0
50.0	39.52	1.19e+0	1.37e+0	1.18e+1	1.44e+0	1.90e+0	1.48e+1	7.28e-1	9.30e-1	1.13e+1	4.40e-1	6.08e-1	2.33e+1
100.0	63.85	2.01e+0	2.31e+0	2.40e+1	2.47e+0	3.23e+0	2.83e+1	1.24e+0	1.54e+0	2.10e+1	6.31e-1	9.37e-1	2.04e+1
150.0	105.39	3.52e+0	4.00e+0	3.87e+1	4.03e+0	5.32e+0	5.11e+1	2.30e+0	2.79e+0	3.21e+1	9.38e-1	1.55e+0	3.99e+1
200.0	105.39	3.52e+0	4.00e+0	3.87e+1	4.03e+0	5.32e+0	5.11e+1	2.30e+0	2.79e+0	3.21e+1	9.38e-1	1.55e+0	3.99e+1
300.0	174.43	6.50e+0	7.33e+0	7.85e+1	6.68e+0	8.74e+0	9.67e+1	4.54e+0	5.44e+0	4.04e+1	1.82e+0	3.25e+0	5.96e+1
400.0	174.43	6.50e+0	7.33e+0	7.85e+1	6.68e+0	8.74e+0	9.67e+1	4.54e+0	5.44e+0	4.04e+1	1.82e+0	3.25e+0	5.96e+1

Table B3. Error values and compression ratios of the ZFP baseline run on the tra_20200101_00_permuted dataset. Note that for some different tolerances the compression ratios are the same, this is probably because of the grouping into “bit planes” that is done by the embedded coding strategy of ZFP.

Tolerance	Ratio	lon-lat			<i>p</i>			TH			PV		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
0.01	2.03	7.18e-4	8.42e-4	3.99e-3	5.11e-4	6.52e-4	3.54e-3	5.11e-4	6.52e-4	3.66e-3	5.11e-4	6.52e-4	3.53e-3
0.05	2.35	2.87e-3	3.37e-3	1.71e-2	2.05e-3	2.61e-3	1.44e-2	2.04e-3	2.61e-3	1.46e-2	2.05e-3	2.61e-3	1.41e-2
0.1	2.55	5.75e-3	6.73e-3	3.29e-2	4.09e-3	5.21e-3	2.79e-2	4.09e-3	5.21e-3	2.74e-2	4.09e-3	5.21e-3	3.00e-2
0.5	3.47	4.34e-2	5.04e-2	2.53e-1	3.27e-2	4.17e-2	2.38e-1	3.27e-2	4.17e-2	2.25e-1	3.12e-2	4.00e-2	2.34e-1
1.0	3.94	8.32e-2	9.64e-2	4.79e-1	6.54e-2	8.34e-2	4.76e-1	6.51e-2	8.30e-2	4.45e-1	5.73e-2	7.41e-2	7.34e-1
2.0	4.52	1.56e-1	1.82e-1	9.53e-1	1.31e-1	1.67e-1	9.22e-1	1.26e-1	1.61e-1	9.29e-1	9.92e-2	1.29e-1	1.50e+0
3.0	4.52	1.56e-1	1.82e-1	9.53e-1	1.31e-1	1.67e-1	9.22e-1	1.26e-1	1.61e-1	9.29e-1	9.92e-2	1.29e-1	1.50e+0
4.0	5.25	2.91e-1	3.41e-1	1.85e+0	2.62e-1	3.34e-1	1.88e+0	2.26e-1	2.89e-1	1.79e+0	1.68e-1	2.17e-1	2.86e+0
5.0	5.25	2.91e-1	3.41e-1	1.85e+0	2.62e-1	3.34e-1	1.88e+0	2.26e-1	2.89e-1	1.79e+0	1.68e-1	2.17e-1	2.86e+0
7.0	5.25	2.91e-1	3.41e-1	1.85e+0	2.62e-1	3.34e-1	1.88e+0	2.26e-1	2.89e-1	1.79e+0	1.68e-1	2.17e-1	2.86e+0
10.0	6.14	5.55e-1	6.54e-1	3.36e+0	5.22e-1	6.65e-1	3.66e+0	3.71e-1	4.74e-1	3.43e+0	2.93e-1	3.71e-1	3.30e+0
15.0	6.14	5.55e-1	6.54e-1	3.36e+0	5.22e-1	6.65e-1	3.66e+0	3.71e-1	4.74e-1	3.43e+0	2.93e-1	3.71e-1	3.30e+0
20.0	7.28	1.06e+0	1.25e+0	7.17e+0	1.02e+0	1.30e+0	7.38e+0	6.15e-1	7.76e-1	6.03e+0	5.27e-1	6.69e-1	6.05e+0
25.0	7.28	1.06e+0	1.25e+0	7.17e+0	1.02e+0	1.30e+0	7.38e+0	6.15e-1	7.76e-1	6.03e+0	5.27e-1	6.69e-1	6.05e+0
50.0	8.78	1.93e+0	2.25e+0	1.31e+1	1.91e+0	2.44e+0	1.49e+1	1.14e+0	1.41e+0	1.04e+1	9.42e-1	1.23e+0	1.13e+1
100.0	10.82	3.46e+0	3.98e+0	2.55e+1	3.42e+0	4.39e+0	2.96e+1	2.21e+0	2.74e+0	2.03e+1	1.73e+0	2.40e+0	2.29e+1
150.0	13.70	6.45e+0	7.36e+0	4.69e+1	5.97e+0	7.69e+0	5.46e+1	4.39e+0	5.43e+0	3.10e+1	2.77e+0	4.30e+0	3.97e+1
200.0	13.70	6.45e+0	7.36e+0	4.69e+1	5.97e+0	7.69e+0	5.46e+1	4.39e+0	5.43e+0	3.10e+1	2.77e+0	4.30e+0	3.97e+1
300.0	18.24	1.26e+1	1.44e+1	9.37e+1	1.03e+1	1.31e+1	9.62e+1	8.76e+0	1.08e+1	4.93e+1	3.61e+0	6.61e+0	5.93e+1
400.0	18.24	1.26e+1	1.44e+1	9.37e+1	1.03e+1	1.31e+1	9.62e+1	8.76e+0	1.08e+1	4.93e+1	3.61e+0	6.61e+0	5.93e+1

Table B4. Error values and compression ratios of the SZ3 baseline run on the tra_20200101_00 dataset.

Tolerance	Ratio	lon-lat			<i>p</i>			TH			PV		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
6.67e-7	2.84	1.13e-4	1.26e-4	2.67e-4	2.86e-4	3.41e-4	6.26e-4	9.31e-5	1.10e-4	1.83e-4	3.70e-5	4.32e-5	7.71e-5
3.33e-6	3.62	5.65e-4	6.32e-4	1.34e-3	1.43e-3	1.70e-3	3.14e-3	4.86e-4	5.67e-4	1.01e-3	1.85e-4	2.16e-4	3.86e-4
6.67e-6	4.09	1.12e-3	1.25e-3	2.68e-3	2.90e-3	3.42e-3	6.28e-3	9.73e-4	1.13e-3	2.01e-3	3.71e-4	4.32e-4	7.71e-4
3.33e-5	5.84	5.60e-3	6.26e-3	1.34e-2	1.45e-2	1.71e-2	3.14e-2	4.86e-3	5.67e-3	1.01e-2	1.85e-3	2.16e-3	3.86e-3
6.67e-5	7.11	1.08e-2	1.22e-2	2.68e-2	2.90e-2	3.43e-2	6.28e-2	9.70e-3	1.13e-2	2.02e-2	3.67e-3	4.29e-3	7.71e-3
1.33e-4	8.93	2.08e-2	2.35e-2	5.36e-2	5.80e-2	6.84e-2	1.26e-1	1.93e-2	2.26e-2	4.04e-2	7.19e-3	8.44e-3	1.54e-2
2.00e-4	10.46	3.00e-2	3.40e-2	8.04e-2	8.67e-2	1.02e-1	1.89e-1	2.88e-2	3.36e-2	6.07e-2	1.05e-2	1.24e-2	2.31e-2
2.67e-4	11.81	3.88e-2	4.42e-2	1.07e-1	1.14e-1	1.36e-1	2.51e-1	3.80e-2	4.45e-2	8.09e-2	1.37e-2	1.63e-2	3.08e-2
3.33e-4	13.11	4.72e-2	5.38e-2	1.34e-1	1.45e-1	1.70e-1	3.14e-1	4.70e-2	5.52e-2	1.01e-1	1.68e-2	2.00e-2	3.86e-2
4.67e-4	15.40	6.30e-2	7.20e-2	1.88e-1	1.97e-1	2.35e-1	4.40e-1	6.44e-2	7.60e-2	1.42e-1	2.26e-2	2.72e-2	5.40e-2
6.67e-4	18.60	8.48e-2	9.73e-2	2.68e-1	2.83e-1	3.34e-1	6.28e-1	8.89e-2	1.06e-1	2.02e-1	3.07e-2	3.73e-2	7.71e-2
1.00e-3	23.31	1.18e-1	1.36e-1	4.01e-1	4.13e-1	4.92e-1	9.43e-1	1.26e-1	1.51e-1	3.03e-1	4.28e-2	5.27e-2	1.16e-1
1.33e-3	28.09	1.48e-1	1.71e-1	5.36e-1	5.38e-1	6.45e-1	1.26e+0	1.58e-1	1.93e-1	4.05e-1	5.39e-2	6.70e-2	1.54e-1
1.67e-3	32.58	1.76e-1	2.03e-1	6.69e-1	6.60e-1	7.94e-1	1.57e+0	1.88e-1	2.31e-1	5.06e-1	6.42e-2	8.04e-2	1.93e-1
3.33e-3	53.23	3.02e-1	3.49e-1	1.34e+0	1.21e+0	1.48e+0	3.14e+0	3.10e-1	3.89e-1	1.01e+0	1.09e-1	1.39e-1	3.86e-1
6.67e-3	90.50	5.36e-1	6.20e-1	2.67e+0	2.13e+0	2.66e+0	6.28e+0	5.05e-1	6.37e-1	2.02e+0	1.85e-1	2.34e-1	7.71e-1
1.00e-2	129.36	7.53e-1	8.62e-1	3.99e+0	2.83e+0	3.62e+0	9.43e+0	6.86e-1	8.58e-1	3.03e+0	2.57e-1	3.21e-1	1.16e+0
1.33e-2	160.90	1.00e+0	1.15e+0	5.30e+0	3.64e+0	4.59e+0	1.26e+1	8.67e-1	1.07e+0	4.05e+0	3.26e-1	4.03e-1	1.54e+0
2.00e-2	223.84	1.44e+0	1.66e+0	7.91e+0	4.99e+0	6.27e+0	1.89e+1	1.24e+0	1.52e+0	6.07e+0	4.46e-1	5.50e-1	2.31e+0
2.67e-2	283.06	1.90e+0	2.18e+0	1.04e+1	6.21e+0	7.80e+0	2.51e+1	1.63e+0	1.97e+0	8.09e+0	5.72e-1	7.03e-1	3.08e+0

Table B5. Error values and compression ratios of the SZ3 baseline run on the tra_20200101_00_permuted dataset.

Tolerance	Ratio	lon-lat			p			TH			PV		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
6.67e-7	1.45	4.50e-5	7.80e-5	2.68e-4	1.24e-4	2.26e-4	6.26e-4	4.52e-5	7.67e-5	1.83e-4	2.75e-5	3.72e-5	7.71e-5
3.33e-6	1.79	3.86e-4	5.05e-4	1.34e-3	1.15e-3	1.54e-3	3.14e-3	4.13e-4	5.24e-4	1.01e-3	1.81e-4	2.14e-4	3.86e-4
6.67e-6	2.11	1.01e-3	1.18e-3	2.68e-3	2.92e-3	3.47e-3	6.28e-3	9.37e-4	1.11e-3	2.01e-3	3.70e-4	4.32e-4	7.71e-4
3.33e-5	2.68	5.79e-3	6.43e-3	1.34e-2	1.51e-2	1.76e-2	3.14e-2	4.86e-3	5.67e-3	1.01e-2	1.85e-3	2.16e-3	3.86e-3
6.67e-5	2.93	1.15e-2	1.28e-2	2.68e-2	3.02e-2	3.52e-2	6.28e-2	9.71e-3	1.13e-2	2.02e-2	3.70e-3	4.32e-3	7.71e-3
1.33e-4	3.22	2.27e-2	2.53e-2	5.36e-2	6.04e-2	7.04e-2	1.26e-1	1.94e-2	2.27e-2	4.04e-2	7.36e-3	8.59e-3	1.54e-2
2.00e-4	3.42	3.37e-2	3.75e-2	8.04e-2	9.06e-2	1.06e-1	1.89e-1	2.91e-2	3.40e-2	6.07e-2	1.10e-2	1.28e-2	2.31e-2
2.67e-4	3.58	4.45e-2	4.97e-2	1.07e-1	1.21e-1	1.41e-1	2.51e-1	3.88e-2	4.52e-2	8.09e-2	1.45e-2	1.70e-2	3.08e-2
3.33e-4	3.72	5.52e-2	6.17e-2	1.34e-1	1.51e-1	1.76e-1	3.14e-1	4.84e-2	5.64e-2	1.01e-1	1.80e-2	2.12e-2	3.86e-2
4.67e-4	3.94	7.64e-2	8.55e-2	1.88e-1	2.11e-1	2.46e-1	4.40e-1	6.74e-2	7.87e-2	1.42e-1	2.50e-2	2.94e-2	5.40e-2
6.67e-4	4.20	1.08e-1	1.21e-1	2.68e-1	3.01e-1	3.51e-1	6.28e-1	9.53e-2	1.12e-1	2.02e-1	3.53e-2	4.16e-2	7.71e-2
1.00e-3	4.55	1.59e-1	1.79e-1	4.02e-1	4.50e-1	5.25e-1	9.43e-1	1.41e-1	1.65e-1	3.03e-1	5.22e-2	6.18e-2	1.16e-1
1.33e-3	4.83	2.10e-1	2.37e-1	5.36e-1	5.97e-1	6.98e-1	1.26e+0	1.85e-1	2.18e-1	4.05e-1	6.89e-2	8.18e-2	1.54e-1
1.67e-3	5.08	2.61e-1	2.95e-1	6.70e-1	7.43e-1	8.70e-1	1.57e+0	2.29e-1	2.70e-1	5.06e-1	8.55e-2	1.02e-1	1.93e-1
3.33e-3	5.95	5.13e-1	5.84e-1	1.34e+0	1.45e+0	1.71e+0	3.14e+0	4.43e-1	5.28e-1	1.01e+0	1.68e-1	2.01e-1	3.86e-1
6.67e-3	7.22	1.02e+0	1.16e+0	2.68e+0	2.82e+0	3.34e+0	6.28e+0	8.68e-1	1.04e+0	2.02e+0	3.29e-1	3.95e-1	7.71e-1
1.00e-2	8.19	1.52e+0	1.74e+0	4.02e+0	4.16e+0	4.95e+0	9.43e+0	1.29e+0	1.56e+0	3.03e+0	4.85e-1	5.84e-1	1.16e+0
1.33e-2	9.01	2.03e+0	2.32e+0	5.36e+0	5.48e+0	6.53e+0	1.26e+1	1.72e+0	2.07e+0	4.05e+0	6.38e-1	7.70e-1	1.54e+0
2.00e-2	10.32	3.05e+0	3.50e+0	8.04e+0	8.11e+0	9.71e+0	1.89e+1	2.59e+0	3.11e+0	6.07e+0	9.21e-1	1.12e+0	2.31e+0
2.67e-2	11.57	4.08e+0	4.67e+0	1.07e+1	1.07e+1	1.29e+1	2.51e+1	3.46e+0	4.16e+0	8.09e+0	1.21e+0	1.48e+0	3.08e+0

Table B6. Error values and compression ratios for psit on the tra_20000101_00 dataset.

Factor	Ratio	lon-lat			p			temperature			potential vorticity			specific humidity		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
1	1.67	5.38e-4	2.82e-3	2.50e+0	4.62e-2	6.88e-2	1.85e-1	2.67e-3	4.20e-3	4.92e-2	5.71e-8	2.98e-7	1.78e-5	3.69e-7	6.98e-7	1.34e-5
5	2.36	9.21e-4	2.94e-3	2.53e+0	3.23e-1	4.82e-1	1.30e+0	3.51e-3	4.97e-3	4.97e-2	6.07e-8	3.08e-7	1.79e-5	4.17e-7	7.85e-7	1.65e-5
10	5.24	1.60e-2	2.34e-2	2.61e+0	6.92e-1	1.03e+0	2.78e+0	5.14e-2	6.99e-2	9.57e-1	4.74e-7	2.20e-6	8.70e-5	5.11e-6	9.89e-6	2.45e-4
15	6.97	3.43e-2	4.99e-2	2.49e+0	1.01e+0	1.51e+0	4.08e+0	1.03e-1	1.40e-1	1.87e+0	9.28e-7	4.32e-6	1.76e-4	1.05e-5	2.04e-5	4.80e-4
20	9.27	6.73e-2	9.72e-2	2.57e+0	1.38e+0	2.06e+0	5.56e+0	1.83e-1	2.50e-1	3.29e+0	1.62e-6	7.54e-6	3.55e-4	1.95e-5	3.75e-5	8.13e-4
25	10.76	9.11e-2	1.31e-1	2.46e+0	1.69e+0	2.53e+0	6.86e+0	2.34e-1	3.21e-1	4.24e+0	2.05e-6	9.57e-6	5.01e-4	2.54e-5	4.90e-5	9.55e-4
30	12.78	1.29e-1	1.86e-1	2.78e+0	2.05e+0	3.08e+0	8.35e+0	3.05e-1	4.21e-1	4.87e+0	2.62e-6	1.23e-5	6.11e-4	3.41e-5	6.56e-5	1.31e-3
40	15.97	1.92e-1	2.73e-1	3.66e+0	2.72e+0	4.08e+0	1.11e+1	4.11e-1	5.77e-1	7.56e+0	3.46e-6	1.66e-5	8.29e-4	4.72e-5	9.11e-5	1.68e-3
50	18.97	2.52e-1	3.57e-1	4.76e+0	3.36e+0	5.07e+0	1.39e+1	5.07e-1	7.20e-1	8.68e+0	4.15e-6	2.03e-5	1.16e-3	5.99e-5	1.16e-4	1.87e-3
75	25.69	3.86e-1	5.47e-1	7.65e+0	4.90e+0	7.46e+0	2.08e+1	7.02e-1	1.02e+0	1.28e+1	5.53e-6	2.82e-5	1.82e-3	8.47e-5	1.65e-4	2.68e-3
100	32.36	5.27e-1	7.56e-1	1.07e+1	6.41e+0	9.83e+0	2.78e+1	8.79e-1	1.31e+0	1.71e+1	6.65e-6	3.50e-5	2.46e-3	1.08e-4	2.14e-4	4.62e-3
150	44.51	7.78e-1	1.14e+0	1.69e+1	9.20e+0	1.43e+1	4.17e+1	1.14e+0	1.74e+0	2.63e+1	8.34e-6	4.60e-5	3.49e-3	1.48e-4	2.97e-4	5.38e-3

Table B7. Error values and compression ratios of the ZFP baseline run on the tra_20000101_00 dataset. Note that for some different tolerances the compression ratios are the same, this is probably because of the grouping into “bit planes” that is done by the embedded coding strategy of ZFP.

Tolerance	Ratio	lon-lat			p			temperature			potential vorticity			specific humidity		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
0.01	4.02	7.42e-4	8.62e-4	4.30e-3	5.30e-4	6.70e-4	3.72e-3	5.30e-4	6.70e-4	3.74e-3	3.15e-5	8.49e-5	5.54e-3	1.39e-4	2.38e-4	5.79e-3
0.05	4.91	2.88e-3	3.33e-3	1.65e-2	2.11e-3	2.67e-3	1.54e-2	2.12e-3	2.68e-3	1.52e-2	8.78e-5	2.59e-4	1.17e-2	3.96e-4	7.14e-4	1.15e-2
0.1	5.49	5.54e-3	6.38e-3	3.53e-2	4.21e-3	5.32e-3	2.92e-2	4.24e-3	5.36e-3	3.02e-2	1.32e-4	4.26e-4	2.16e-2	6.31e-4	1.20e-3	2.29e-2
0.5	8.20	3.64e-2	4.23e-2	2.21e-1	3.21e-2	4.09e-2	2.35e-1	3.35e-2	4.23e-2	2.39e-1	2.12e-4	9.17e-4	3.91e-2	2.14e-3	4.51e-3	3.12e-2
1.0	9.64	6.71e-2	7.83e-2	4.37e-1	6.22e-2	7.97e-2	4.86e-1	6.52e-2	8.23e-2	4.69e-1	2.11e-4	8.97e-4	4.44e-2	2.11e-3	4.39e-3	2.25e-2
2.0	11.53	1.24e-1	1.45e-1	8.44e-1	1.19e-1	1.54e-1	9.36e-1	1.23e-1	1.55e-1	9.53e-1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
3.0	11.53	1.24e-1	1.45e-1	8.44e-1	1.19e-1	1.54e-1	9.36e-1	1.23e-1	1.55e-1	9.53e-1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
4.0	14.11	2.31e-1	2.70e-1	1.67e+0	2.25e-1	2.96e-1	1.85e+0	2.21e-1	2.81e-1	1.82e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
5.0	14.11	2.31e-1	2.70e-1	1.67e+0	2.25e-1	2.96e-1	1.85e+0	2.21e-1	2.81e-1	1.82e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
7.0	14.11	2.31e-1	2.70e-1	1.67e+0	2.25e-1	2.96e-1	1.85e+0	2.21e-1	2.81e-1	1.82e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
10.0	17.77	4.27e-1	4.97e-1	3.20e+0	4.21e-1	5.60e-1	3.79e+0	3.79e-1	4.87e-1	3.66e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
15.0	17.77	4.27e-1	4.97e-1	3.20e+0	4.21e-1	5.60e-1	3.79e+0	3.79e-1	4.87e-1	3.66e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
20.0	23.13	7.65e-1	8.87e-1	6.24e+0	7.66e-1	1.04e+0	7.39e+0	6.29e-1	8.13e-1	6.80e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
25.0	23.13	7.65e-1	8.87e-1	6.24e+0	7.66e-1	1.04e+0	7.39e+0	6.29e-1	8.13e-1	6.80e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
50.0	31.07	1.40e+0	1.61e+0	1.23e+1	1.40e+0	1.88e+0	1.46e+1	1.04e+0	1.34e+0	1.28e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
100.0	43.36	2.44e+0	2.84e+0	2.76e+1	2.48e+0	3.32e+0	2.76e+1	1.76e+0	2.26e+0	2.26e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
150.0	62.83	4.26e+0	4.98e+0	4.76e+1	4.28e+0	5.72e+0	5.74e+1	3.05e+0	3.94e+0	4.23e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
200.0	62.83	4.26e+0	4.98e+0	4.76e+1	4.28e+0	5.72e+0	5.74e+1	3.05e+0	3.94e+0	4.23e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
300.0	93.96	7.59e+0	8.82e+0	9.97e+1	7.32e+0	9.77e+0	1.11e+2	5.41e+0	6.94e+0	6.84e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
400.0	93.96	7.59e+0	8.82e+0	9.97e+1	7.32e+0	9.77e+0	1.11e+2	5.41e+0	6.94e+0	6.84e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2

Table B8. Error values and compression ratios of the ZFP baseline run on the `tra_20000101_00_permuted` dataset. Note that for some different tolerances the compression ratios are the same, this is probably because of the grouping into “bit planes” that is done by the embedded coding strategy of ZFP.

Tolerance	Ratio	lon-lat			p			temperature			potential vorticity			specific humidity		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
0.01	3.09	7.44e-4	8.65e-4	4.32e-3	5.30e-4	6.70e-4	3.75e-3	5.30e-4	6.70e-4	3.81e-3	1.11e-4	2.16e-4	5.78e-3	2.51e-4	3.33e-4	3.24e-3
0.05	3.65	2.95e-3	3.43e-3	1.68e-2	2.12e-3	2.68e-3	1.50e-2	2.12e-3	2.68e-3	1.52e-2	2.14e-4	5.74e-4	2.30e-2	8.43e-4	1.18e-3	2.29e-2
0.1	4.00	5.77e-3	6.68e-3	3.32e-2	4.24e-3	5.36e-3	3.03e-2	4.24e-3	5.36e-3	3.07e-2	2.37e-4	7.96e-4	4.40e-2	1.44e-3	2.17e-3	4.62e-2
0.5	5.49	4.00e-2	4.64e-2	2.65e-1	3.39e-2	4.29e-2	2.44e-1	3.39e-2	4.29e-2	2.37e-1	2.11e-4	8.97e-4	4.69e-2	2.11e-3	4.41e-3	3.12e-2
1.0	6.20	7.39e-2	8.65e-2	4.89e-1	6.77e-2	8.56e-2	4.78e-1	6.74e-2	8.52e-2	4.70e-1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
2.0	7.08	1.42e-1	1.67e-1	9.30e-1	1.35e-1	1.70e-1	9.69e-1	1.32e-1	1.66e-1	9.54e-1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
3.0	7.08	1.42e-1	1.67e-1	9.30e-1	1.35e-1	1.70e-1	9.69e-1	1.32e-1	1.66e-1	9.54e-1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
4.0	8.23	2.76e-1	3.25e-1	1.78e+0	2.65e-1	3.35e-1	1.99e+0	2.45e-1	3.11e-1	1.84e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
5.0	8.23	2.76e-1	3.25e-1	1.78e+0	2.65e-1	3.35e-1	1.99e+0	2.45e-1	3.11e-1	1.84e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
7.0	8.23	2.76e-1	3.25e-1	1.78e+0	2.65e-1	3.35e-1	1.99e+0	2.45e-1	3.11e-1	1.84e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
10.0	9.73	5.20e-1	6.09e-1	3.44e+0	5.08e-1	6.44e-1	3.82e+0	4.35e-1	5.56e-1	3.79e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
15.0	9.73	5.20e-1	6.09e-1	3.44e+0	5.08e-1	6.44e-1	3.82e+0	4.35e-1	5.56e-1	3.79e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
20.0	11.61	9.53e-1	1.11e+0	6.80e+0	9.46e-1	1.20e+0	7.63e+0	7.31e-1	9.38e-1	6.93e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
25.0	11.61	9.53e-1	1.11e+0	6.80e+0	9.46e-1	1.20e+0	7.63e+0	7.31e-1	9.38e-1	6.93e+0	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
50.0	13.96	1.73e+0	1.99e+0	1.34e+1	1.71e+0	2.20e+0	1.49e+1	1.23e+0	1.55e+0	1.43e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
100.0	16.88	3.25e+0	3.71e+0	2.83e+1	3.02e+0	3.90e+0	2.89e+1	2.27e+0	2.81e+0	2.59e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
150.0	20.87	6.33e+0	7.23e+0	4.92e+1	5.24e+0	6.71e+0	5.42e+1	4.42e+0	5.47e+0	4.22e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
200.0	20.87	6.33e+0	7.23e+0	4.92e+1	5.24e+0	6.71e+0	5.42e+1	4.42e+0	5.47e+0	4.22e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
300.0	27.05	1.25e+1	1.43e+1	9.40e+1	9.43e+0	1.18e+1	1.11e+2	8.77e+0	1.09e+1	6.98e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2
400.0	27.05	1.25e+1	1.43e+1	9.40e+1	9.43e+0	1.18e+1	1.11e+2	8.77e+0	1.09e+1	6.98e+1	2.11e-4	8.97e-4	3.91e-2	2.11e-3	4.39e-3	2.25e-2

Table B9. Error values and compression ratios of the SZ3 baseline run on the `tra_20000101_00` dataset.

Tolerance	Ratio	lon-lat			p			temperature			potential vorticity			specific humidity		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
6.67e-7	3.27	1.60e-4	1.81e-4	3.79e-4	4.98e-4	5.76e-4	1.00e-3	3.61e-5	4.41e-5	7.63e-5	2.22e-8	2.56e-8	4.44e-8	7.47e-9	8.64e-9	1.50e-8
3.33e-6	4.33	8.06e-4	9.09e-4	1.90e-3	2.49e-3	2.88e-3	5.00e-3	2.17e-4	2.52e-4	4.27e-4	1.10e-7	1.27e-7	2.22e-7	3.74e-8	4.32e-8	7.49e-8
6.67e-6	4.94	1.62e-3	1.82e-3	3.79e-3	4.99e-3	5.76e-3	1.00e-2	4.34e-4	5.03e-4	8.70e-4	1.53e-7	1.92e-7	4.44e-7	7.48e-8	8.64e-8	1.50e-7
3.33e-5	6.96	8.12e-3	9.15e-3	1.90e-2	2.50e-2	2.89e-2	5.00e-2	2.20e-3	2.54e-3	4.39e-3	5.46e-7	7.42e-7	2.22e-6	3.73e-7	4.31e-7	7.49e-7
6.67e-5	8.15	1.62e-2	1.83e-2	3.79e-2	5.00e-2	5.77e-2	1.00e-1	4.40e-3	5.08e-3	8.79e-3	9.13e-7	1.32e-6	4.44e-6	7.42e-7	8.58e-7	1.50e-6
1.33e-4	9.71	3.25e-2	3.66e-2	7.59e-2	1.00e-1	1.16e-1	2.00e-1	8.80e-3	1.02e-2	1.76e-2	1.49e-6	2.36e-6	8.88e-6	1.49e-6	1.72e-6	3.00e-6
2.00e-4	10.83	4.87e-2	5.49e-2	1.14e-1	1.50e-1	1.73e-1	3.00e-1	1.32e-2	1.52e-2	2.64e-2	2.11e-6	3.41e-6	1.33e-5	2.23e-6	2.58e-6	4.49e-6
2.67e-4	11.86	6.50e-2	7.32e-2	1.52e-1	2.00e-1	2.31e-1	4.00e-1	1.76e-2	2.03e-2	3.52e-2	2.47e-6	4.28e-6	1.78e-5	3.02e-6	3.47e-6	5.99e-6
3.33e-4	12.68	8.12e-2	9.15e-2	1.90e-1	2.49e-1	2.88e-1	5.00e-1	2.20e-2	2.54e-2	4.40e-2	2.92e-6	5.16e-6	2.22e-5	3.77e-6	4.33e-6	7.49e-6
4.67e-4	13.77	9.14e-2	1.05e-1	2.66e-1	3.50e-1	4.04e-1	7.00e-1	3.08e-2	3.55e-2	6.16e-2	3.57e-6	6.71e-6	3.11e-5	5.19e-6	5.99e-6	1.05e-5
6.67e-4	15.94	1.62e-1	1.83e-1	3.79e-1	4.99e-1	5.77e-1	1.00e+0	4.40e-2	5.08e-2	8.80e-2	4.72e-6	9.09e-6	4.44e-5	7.50e-6	8.69e-6	1.50e-5
1.00e-3	18.39	1.90e-1	2.18e-1	5.69e-1	7.50e-1	8.66e-1	1.50e+0	6.60e-2	7.62e-2	1.32e-1	5.94e-6	1.23e-5	6.66e-5	1.17e-5	1.33e-5	2.25e-5
1.33e-3	20.50	2.25e-1	2.70e-1	7.59e-1	1.00e+0	1.16e+0	2.00e+0	8.80e-2	1.02e-1	1.76e-1	7.23e-6	1.53e-5	8.88e-5	1.49e-5	1.72e-5	3.00e-5
1.67e-3	22.64	2.76e-1	3.32e-1	9.48e-1	1.25e+0	1.44e+0	2.50e+0	1.10e-1	1.27e-1	2.20e-1	8.71e-6	1.88e-5	1.11e-4	1.93e-5	2.25e-5	3.74e-5
3.33e-3	31.23	5.17e-1	6.29e-1	1.90e+0	2.49e+0	2.88e+0	5.00e+0	2.20e-1	2.54e-1	4.40e-1	1.41e-5	3.19e-5	2.22e-4	3.77e-5	4.44e-5	7.49e-5
6.67e-3	45.33	9.59e-1	1.18e+0	3.79e+0	4.99e+0	5.79e+0	1.00e+1	4.40e-1	5.08e-1	8.80e-1	2.03e-5	5.05e-5	4.44e-4	7.51e-5	9.09e-5	1.50e-4
1.00e-2	53.45	2.44e+0	2.75e+0	5.69e+0	6.87e+0	8.19e+0	1.50e+1	6.60e-1	7.62e-1	1.32e+0	2.79e-5	7.12e-5	6.66e-4	1.15e-4	1.43e-4	2.25e-4
1.33e-2	66.81	3.02e+0	3.50e+0	7.59e+0	1.01e+1	1.16e+1	2.00e+1	8.80e-1	1.02e+0	1.76e+0	5.41e-5	1.21e-4	8.88e-4	1.49e-4	1.74e-4	3.00e-4
2.00e-2	80.59	4.51e+0	5.24e+0	1.14e+1	1.57e+1	1.80e+1	3.00e+1	1.32e+0	1.52e+0	2.64e+0	8.10e-5	1.82e-4	1.33e-3	2.17e-4	2.48e-4	4.49e-4
2.67e-2	89.20	3.27e+0	4.02e+0	1.52e+1	1.99e+1	2.28e+1	4.00e+1	1.76e+0	2.03e+0	3.52e+0	1.12e-4	2.49e-4	1.78e-3	1.68e-4	2.11e-4	5.99e-4

Table B10. Error values and compression ratios of the SZ3 baseline run on the tra_20000101_00_permuted dataset.

Tolerance	Ratio	lon-lat			ρ			temperature			potential vorticity			specific humidity		
		L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity	L1	RMSE	L-infinity
6.67e-7	2.71	1.58e-4	1.80e-4	3.79e-4	4.95e-4	5.75e-4	1.00e-3	3.59e-5	4.40e-5	7.63e-5	2.22e-8	2.56e-8	4.44e-8	7.43e-9	8.61e-9	1.50e-8
3.33e-6	3.42	8.00e-4	9.05e-4	1.90e-3	2.48e-3	2.87e-3	5.00e-3	2.16e-4	2.51e-4	4.27e-4	1.11e-7	1.28e-7	2.22e-7	3.74e-8	4.32e-8	7.49e-8
6.67e-6	3.80	1.61e-3	1.82e-3	3.79e-3	4.98e-3	5.76e-3	1.00e-2	4.34e-4	5.03e-4	8.70e-4	2.22e-7	2.56e-7	4.44e-7	7.48e-8	8.64e-8	1.50e-7
3.33e-5	5.14	8.12e-3	9.15e-3	1.90e-2	2.50e-2	2.89e-2	5.00e-2	2.20e-3	2.54e-3	4.39e-3	1.10e-6	1.27e-6	2.22e-6	3.74e-7	4.32e-7	7.49e-7
6.67e-5	5.96	1.62e-2	1.83e-2	3.79e-2	5.00e-2	5.78e-2	1.00e-1	4.40e-3	5.08e-3	8.79e-3	2.13e-6	2.50e-6	4.44e-6	7.49e-7	8.65e-7	1.50e-6
1.33e-4	6.99	3.25e-2	3.66e-2	7.59e-2	1.00e-1	1.16e-1	2.00e-1	8.80e-3	1.02e-2	1.76e-2	4.28e-6	5.02e-6	8.88e-6	1.50e-6	1.73e-6	3.00e-6
2.00e-4	7.69	4.87e-2	5.49e-2	1.14e-1	1.50e-1	1.73e-1	3.00e-1	1.32e-2	1.52e-2	2.64e-2	6.50e-6	7.76e-6	1.33e-5	2.25e-6	2.59e-6	4.49e-6
2.67e-4	8.25	6.50e-2	7.32e-2	1.52e-1	2.00e-1	2.31e-1	4.00e-1	1.76e-2	2.03e-2	3.52e-2	8.40e-6	1.03e-5	1.78e-5	3.00e-6	3.46e-6	5.99e-6
3.33e-4	8.70	8.12e-2	9.15e-2	1.90e-1	2.50e-1	2.88e-1	5.00e-1	2.20e-2	2.54e-2	4.40e-2	1.06e-5	1.35e-5	2.22e-5	3.75e-6	4.33e-6	7.49e-6
4.67e-4	9.42	1.14e-1	1.28e-1	2.66e-1	3.50e-1	4.04e-1	7.00e-1	3.08e-2	3.55e-2	6.16e-2	1.54e-5	1.99e-5	3.11e-5	5.24e-6	6.05e-6	1.05e-5
6.67e-4	10.58	1.62e-1	1.83e-1	3.79e-1	5.00e-1	5.77e-1	1.00e+0	4.40e-2	5.08e-2	8.80e-2	2.26e-5	2.84e-5	4.44e-5	7.48e-6	8.65e-6	1.50e-5
1.00e-3	12.41	2.44e-1	2.75e-1	5.69e-1	7.50e-1	8.66e-1	1.50e+0	6.60e-2	7.62e-2	1.32e-1	3.72e-5	4.63e-5	6.66e-5	1.12e-5	1.29e-5	2.25e-5
1.33e-3	13.87	3.25e-1	3.66e-1	7.59e-1	1.00e+0	1.16e+0	2.00e+0	8.80e-2	1.02e-1	1.76e-1	3.96e-5	5.32e-5	8.88e-5	1.49e-5	1.74e-5	3.00e-5
1.67e-3	14.71	4.06e-1	4.58e-1	9.49e-1	1.25e+0	1.44e+0	2.50e+0	1.10e-1	1.27e-1	2.20e-1	3.24e-5	4.29e-5	1.11e-4	1.87e-5	2.18e-5	3.74e-5
3.33e-3	20.20	8.12e-1	9.15e-1	1.90e+0	2.49e+0	2.88e+0	5.00e+0	2.20e-1	2.54e-1	4.40e-1	6.13e-5	8.34e-5	2.22e-4	3.73e-5	4.32e-5	7.49e-5
6.67e-3	27.86	1.62e+0	1.83e+0	3.79e+0	5.00e+0	5.81e+0	1.00e+1	4.40e-1	5.08e-1	8.80e-1	1.18e-4	1.62e-4	4.44e-4	7.46e-5	8.71e-5	1.50e-4
1.00e-2	36.51	2.44e+0	2.74e+0	5.69e+0	6.94e+0	8.22e+0	1.50e+1	6.60e-1	7.62e-1	1.32e+0	1.67e-4	2.33e-4	6.66e-4	1.12e-4	1.35e-4	2.25e-4
1.33e-2	41.10	3.25e+0	3.66e+0	7.59e+0	1.00e+1	1.16e+1	2.00e+1	8.80e-1	1.02e+0	1.76e+0	2.16e-4	3.02e-4	8.88e-4	1.49e-4	1.76e-4	3.00e-4
2.00e-2	47.07	4.87e+0	5.49e+0	1.14e+1	1.58e+1	1.82e+1	3.00e+1	1.32e+0	1.52e+0	2.64e+0	2.77e-4	4.14e-4	1.33e-3	2.16e-4	2.69e-4	4.49e-4
2.67e-2	50.70	6.49e+0	7.32e+0	1.52e+1	1.99e+1	2.30e+1	4.00e+1	1.76e+0	2.03e+0	3.52e+0	3.38e-4	5.23e-4	1.78e-3	2.81e-4	3.58e-4	5.99e-4

Table B11. Jensen–Shannon divergence metrics for different time steps for trajectories starting at 1 hPa.

Dataset	2D			Pressure		
	72 h	120 h	168 h	72 h	120 h	168 h
Noise 0.01	9.94e-4	1.73e-3	2.53e-3	3.10e-4	8.43e-3	3.28e-4
Noise 0.05	4.46e-3	5.87e-3	7.03e-3	1.07e-1	1.10e-1	1.19e-2
Noise 0.1	7.15e-3	9.38e-3	1.03e-2	2.81e-1	1.09e-1	4.97e-2
Compressed 1.67x	4.27e-6	3.78e-6	4.27e-6	1.99e-7	9.15e-8	1.67e-7
Compressed 12.8x	4.17e-4	1.01e-3	1.98e-3	1.18e-4	1.90e-4	1.27e-4
Compressed 44.5x	2.95e-3	7.88e-3	1.25e-2	6.94e-4	1.50e-3	1.40e-3

Table B12. Jensen–Shannon divergence metrics for different time steps for trajectories starting at 500 hPa.

Dataset	2D			Pressure		
	72 h	120 h	168 h	72 h	120 h	168 h
Noise 0.01	5.51e-5	2.84e-4	9.68e-4	4.56e-2	3.36e-1	1.03e-1
Noise 0.05	2.43e-4	1.02e-3	2.33e-3	4.59e-2	3.36e-1	1.01e-1
Noise 0.1	4.42e-4	1.59e-3	3.34e-3	4.51e-2	3.28e-1	9.66e-2
Compressed 1.67x	1.93e-6	1.31e-6	1.10e-6	4.56e-2	3.35e-1	1.03e-1
Compressed 12.8x	9.03e-4	1.45e-3	1.79e-3	4.27e-2	3.36e-1	9.93e-2
Compressed 44.5x	6.76e-3	1.01e-2	1.23e-2	3.27e-2	3.16e-1	8.84e-2

Table B13. Jensen–Shannon divergence metrics for different time steps for trajectories starting at 1000 hPa.

Dataset	2D			Pressure		
	72 h	120 h	168 h	72 h	120 h	168 h
Noise 0.01	7.80e-5	5.40e-4	1.64e-3	3.80e-1	1.73e-1	2.06e-1
Noise 0.05	3.45e-4	1.51e-3	3.23e-3	3.74e-1	1.67e-1	2.05e-1
Noise 0.1	5.82e-4	2.21e-3	4.24e-3	3.74e-1	1.72e-1	2.06e-1
Compressed 1.67x	2.03e-6	1.64e-6	2.33e-6	3.80e-1	1.73e-1	2.06e-1
Compressed 12.8x	7.36e-4	1.15e-3	1.61e-3	3.79e-1	1.67e-1	2.04e-1
Compressed 44.5x	5.65e-3	1.01e-2	1.28e-2	3.62e-1	1.54e-1	1.96e-1

Appendix C: User Manual

In this section a small user manual for psit is given. For a detailed description please refer to (Pietak, 2025a). The psit compression pipeline has been implemented in Python, as this allowed for fast development and quick iteration of different ideas. The entire compression and decompression logic has been implemented in a class called `Psit`, which exposes two functions: `compress()` and `decompress()`. The `compress()` function takes an `xarray.Dataset` (with additional configuration parameters) as input and compresses it to a file on disk. The `decompress()` function takes the location of such a compressed file and returns an `xarray.Dataset` with the decompressed trajectory data. The additional configuration parameters, which can be passed to the `compress()` function, and a brief description are provided below.

- `dataset`. The `xarray.Dataset` we wish to compress.
- `filename`. The file in which the compressed data should be stored.
- `crf`. The compression factor.
- `exclude`. List of data variables which should not be compressed.
- `method`. The compression algorithm that should be used.
- `color_method`. The color encoding method that should be used.
- `color_bits`. The number of bits to use per color channel (when using color encoding).
- `delta_method`. The delta encoding method that should be used.
- `local`. If the trajectory starting positions are in a local domain and not global.
- `bin]` The number of pressure bins that should be used.
- `mapping_func`. Which mapping method should be used.
- `num_workers`. The number of parallel workers that should be used.
- `factor`. The ratio between number of grid points and trajectories.
- `lon_name`. Name of the longitude variable in the dataset.
- `lat_name`. Name of the latitude variable in the dataset.

- `p_name`. Name of the pressure variable in the dataset.

Note that some of these parameters, like `crf` or `method`, allow the individual setting of the parameter for each of the data variables available (using a dictionary mapping from variable name to value) in the dataset.

In addition to the `Psit` class, a CLI wrapper has been written called `psitcli`. This CLI wrapper allows for the easy compression of netCDF (Unidata, 2024) files from the command line. The program can be run with the command `python psitcli.py {compress/decompress/test} [options] in_file out_file`. The first option that needs to be defined is the mode to be used, this can either be `compression`, `decompression`, or `test`. The first two are used for compression and decompression, while the `test` mode will perform a compression followed by a decompression and the calculation of performance metrics such as the L1 error that was created. These additional metrics are stored as a Python pickle dump in a file called `<out_file>_data.pick`. The `compress` and `test` modes both take additional options, which mirror the ones accepted by the `compress()` function of the `Psit` class. Alternatively, these options can be defined via a YAML configuration file, which is then passed to `psitcli`. An example of such a configuration file is given in Listing 1.

Code availability. The version of psit referenced in this paper is available at <https://doi.org/10.5281/zenodo.14888491> (Pietak, 2025a). An up to date version can be found on GitHub <https://github.com/apietak/psit> (Pietak, 2025b).

Data availability. Instructions on how the data used in this paper can be obtained is available at <https://doi.org/10.5281/zenodo.19887527> (Pietak, 2026).

Author contributions. AP, LH, and LF designed psit. AP implemented and ran the experiments using psit. MS and SS provided domain specific knowledge and general advice. TH supervised the project. AP prepared the manuscript with the help of all other co-authors.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. The authors bear the ultimate responsibility for providing appropriate place names. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Acknowledgements. We would like to thank CSCS the Swiss super-computer center for providing us with hardware resources without which this project could not have been done.

AP would like to thank James Murdoch MacGregor (J. T. McIntosh), for writing the novelet “Humanoid Sacrifice” (McIntosh, 1964) and creating the name “psit”.

Review statement. This paper was edited by David Ham and reviewed by Robert Underwood and one anonymous referee.

References

- Al-Salamee, B. A. and Al-Shammary, D.: Survey Analysis for Medical Image Compression Techniques, in: *Communication and Intelligent Systems*, edited by: Sharma, H., Gupta, M. K., Tomar, G. S., and Lipo, W., Springer Singapore, Singapore, 241–264, ISBN 978-981-16-1089-9, 2021.
- Bakels, L., Blaschek, M., Dütsch, M., Plach, A., Lechner, V., Brack, G., Haimberger, L., and Stohl, A.: LARA: a Lagrangian Reanalysis based on ERA5 spanning from 1940 to 2023, *Earth Syst. Sci. Data*, 17, 4569–4585, <https://doi.org/10.5194/essd-17-4569-2025>, 2025.
- Baker, A. H., Xu, H., Dennis, J. M., Levy, M. N., Nychka, D., Mickelson, S. A., Edwards, J., Vertenstein, M., and Wegener, A.: A methodology for evaluating the impact of data compression on climate simulation data, in: *Proceedings of the 23rd International Symposium on High-Performance Parallel and Distributed Computing, HPDC '14*, Association for Computing Machinery, New York, NY, USA, 203–214, ISBN 9781450327497, <https://doi.org/10.1145/2600212.2600217>, 2014.
- Baker, A. H., Hammerling, D. M., Mickelson, S. A., Xu, H., Stolpe, M. B., Naveau, P., Sanderson, B., Ebert-Uphoff, I., Samarasinghe, S., De Simone, F., Carbone, F., Gencarelli, C. N., Dennis, J. M., Kay, J. E., and Lindstrom, P.: Evaluating lossy data compression on climate simulation data within a large ensemble, *Geosci. Model Dev.*, 9, 4381–4403, <https://doi.org/10.5194/gmd-9-4381-2016>, 2016.
- Ballester-Ripoll, R. and Pajarola, R.: Lossy volume compression using Tucker truncation and thresholding, *The Visual Computer*, 32, 1433–1446, <https://doi.org/10.1007/s00371-015-1130-y>, 2015.
- Ballester-Ripoll, R., Lindstrom, P., and Pajarola, R.: TTHRESH: Tensor Compression for Multidimensional Visual Data, *IEEE T. Vis. Comput. Gr.*, 26, 2891–2903, 2019.
- Bertsimas, D. and Tsitsiklis, J.: *Introduction to Linear Optimization*, Athena Scientific, ISBN 1886529191, 1997.
- Browning, K. A., Hardman, M. E., Harrold, T. W., and Pardee, C. W.: The structure of rainbands within a mid-latitude depression, *Q. J. Roy. Meteor. Soc.*, 99, 215–231, <https://doi.org/10.1002/qj.49709942002>, 1973.
- Burkard, R., Dell’Amico, M., and Martello, S.: *Assignment Problems*, Society for Industrial and Applied Mathematics, <https://doi.org/10.1137/1.9781611972238>, 2012.
- Calhoun, D. A., Helzel, C., and LeVeque, R. J.: Logically Rectangular Grids and Finite Volume Methods for PDEs in Circular and Spherical Domains, *SIAM Rev.*, 50, 723–752, <https://doi.org/10.1137/060664094>, 2008.
- Chiarot, G. and Silvestri, C.: Time Series Compression Survey, *ACM Comput. Surv.*, 55, <https://doi.org/10.1145/3560814>, 2023.
- Chino, M., Nakayama, H., Nagai, H., Terada, H., Katata, G., and Yamazawa, H.: Preliminary estimation of release amounts of ^{131}I and ^{137}Cs accidentally discharged from the Fukushima Daiichi nuclear power plant into the atmosphere, *J. Nucl. Sci. Technol.*, 48, 1129–1134, <https://doi.org/10.1080/18811248.2011.9711799>, 2011.
- Dey, D., Aldama Campino, A., and Döös, K.: Atmospheric water transport connectivity within and between ocean basins and land, *Hydrol. Earth Syst. Sci.*, 27, 481–493, <https://doi.org/10.5194/hess-27-481-2023>, 2023.
- Di, S., Liu, J., Zhao, K., Liang, X., Underwood, R., Zhang, Z., Shah, M., Huang, Y., Huang, J., Yu, X., Ren, C., Guo, H., Wilkins, G., Tao, D., Tian, J., Jin, S., Jian, Z., Wang, D., Rahman, M. H., Zhang, B., Song, S., Calhoun, J., Li, G., Yoshii, K., Alharthi, K., and Cappello, F.: A Survey on Error-Bounded Lossy Compression for Scientific Datasets, *ACM Comput. Surv.*, 57, <https://doi.org/10.1145/3733104>, 2025.
- Diffenderfer, J., Fox, A., Hittinger, J., Sanders, G., and Lindstrom, P.: Error Analysis of ZFP Compression for Floating-Point Data, *SIAM J. Sci. Comput.*, 41, A1867–A1898, <https://doi.org/10.1137/18M1168832>, 2019.
- Grams, C. M., Wernli, H., Böttcher, M., Čampa, J., Corsmeier, U., Jones, S. C., Keller, J. H., Lenz, C.-J., and Wiegand, L.: The key role of diabatic processes in modifying the upper-tropospheric wave guide: a North Atlantic case-study, *Q. J. Roy. Meteor. Soc.*, 137, 2174–2193, <https://doi.org/10.1002/qj.891>, 2011.
- Han, T., Chen, Z., Guo, S., Xu, W., and Bai, L.: CRA5: Extreme Compression of ERA5 for Portable Global Climate and Weather Research via an Efficient Variational Transformer, *arXiv [preprint]*, <https://doi.org/10.48550/arxiv.2405.03376>, 8 May 2024.
- Hoefler, T., Stevens, B., Prein, A. F., Baehr, J., Schulthess, T., Stocker, T. F., Taylor, J., Klocke, D., Manninen, P., Forster, P. M., Kolling, T., Gruber, N., Anzt, H., Frauen, C., Ziemer, F., Klawer, M., Kashinath, K., Schar, C., Fuhrer, O., and Lawrence, B. N.: Earth Virtualization Engines: A Technical Perspective, *Comput. Sci. Eng.*, 25, 50–59, <https://doi.org/10.1109/MCSE.2023.3311148>, 2023.
- Hoefler, T., Calotou, A., Dipankar, A., Schulthess, T., Lapillonne, X., and Fuhrer, O.: Towards Specialized Supercomputers for Climate Sciences: Computational Requirements of the Icosahedral Nonhydrostatic Weather and Climate Model, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2405.13043>, 18 May 2024.
- Holton, J. R., Haynes, P. H., McIntyre, M. E., Douglass, A. R., Rood, R. B., and Pfister, L.: Stratosphere-troposphere exchange, *Rev. Geophys.*, 33, 403–439, <https://doi.org/10.1029/95RG02097>, 1995.
- IAEA: The Fukushima Daiichi Accident, Non-serial Publications, International Atomic Energy Agency, Vienna, ISBN 978-92-0-107015-9, <https://www.iaea.org/publications/10962/the-fukushima-daiichi-accident> (last access: 9 April 2026), 2015.

- Joos, H. and Wernli, H.: Influence of microphysical processes on the potential vorticity development in a warm conveyor belt: a case-study with the limited-area model COSMO, *Q. J. Roy. Meteor. Soc.*, 138, 407–418, <https://doi.org/10.1002/qj.934>, 2012.
- Katata, G., Ota, M., Terada, H., Chino, M., and Nagai, H.: Atmospheric discharge and dispersion of radionuclides during the Fukushima Dai-ichi Nuclear Power Plant accident. Part I: Source term estimation and local-scale atmospheric dispersion in early phase of the accident, *J. Environ. Radioactiv.*, 109, 103–113, <https://doi.org/10.1016/j.jenvrad.2012.02.006>, 2012.
- Keune, J., Schumacher, D. L., and Miralles, D. G.: A unified framework to estimate the origins of atmospheric moisture and heat using Lagrangian models, *Geosci. Model Dev.*, 15, 1875–1898, <https://doi.org/10.5194/gmd-15-1875-2022>, 2022.
- Liang, X., Di, S., Tao, D., Li, S., Li, S., Guo, H., Chen, Z., and Cappello, F.: Error-Controlled Lossy Compression Optimized for High Compression Ratios of Scientific Datasets, in: 2018 IEEE International Conference on Big Data (Big Data), 438–447, <https://doi.org/10.1109/BigData.2018.8622520>, 2018.
- Liang, X., Di, S., Cappello, F., Raj, M., Liu, C., Ono, K., Chen, Z., Peterka, T., and Guo, H.: Toward Feature-Preserving Vector Field Compression, *IEEE T. Vis. Comput. Gr.*, 29, 5434–5450, <https://doi.org/10.1109/TVCG.2022.3214821>, 2023a.
- Liang, X., Zhao, K., Di, S., Li, S., Underwood, R., Gok, A. M., Tian, J., Deng, J., Calhoun, J. C., Tao, D., Chen, Z., and Cappello, F.: SZ3: A Modular Framework for Composing Prediction-Based Error-Bounded Lossy Compressors, *IEEE Transactions on Big Data*, 9, 485–498, <https://doi.org/10.1109/TBDATA.2022.3201176>, 2023b.
- Lindstrom, P.: Fixed-Rate Compressed Floating-Point Arrays, *IEEE T. Vis. Comput. Gr.*, 20, <https://doi.org/10.1109/TVCG.2014.2346458>, 2014.
- Madonna, E., Wernli, H., Joos, H., and Martius, O.: Warm Conveyor Belts in the ERA-Interim Dataset (1979–2010). Part I: Climatology and Potential Vorticity Evolution, *J. Climate*, 27, 3–26, <https://doi.org/10.1175/JCLI-D-12-00720.1>, 2014.
- Makris, A., Kontopoulos, I., Alimisis, P., and Tserpes, K.: A Comparison of Trajectory Compression Algorithms Over AIS Data, *IEEE Access*, 9, 92516–92530, <https://doi.org/10.1109/ACCESS.2021.3092948>, 2021.
- McIntosh, J. T.: Humanoid Sacrifice, *The Magazine of Fantasy and Science Fiction*, https://archive.org/details/Fantasy_Science_Fiction_v026n03_1964-03_PDF/page/n127/mode/2up (last access: 13 April 2026), 1964.
- Pérez-Muñuzuri, V., Eiras-Barca, J., and Garaboa-Paz, D.: Tagging moisture sources with Lagrangian and inertial tracers: application to intense atmospheric river events, *Earth Syst. Dynam.*, 9, 785–795, <https://doi.org/10.5194/esd-9-785-2018>, 2018.
- Pfahl, S., Madonna, E., Boettcher, M., Joos, H., and Wernli, H.: Warm Conveyor Belts in the ERA-Interim Dataset (1979–2010). Part II: Moisture Origin and Relevance for Precipitation, *J. Climate*, 27, 27–40, <https://doi.org/10.1175/JCLI-D-13-00223.1>, 2014.
- Pietak, A.: `apietak/psit: psit v1.0.0`, Zenodo [code], <https://doi.org/10.5281/zenodo.14888491>, 2025a.
- Pietak, A.: `psit`, GitHub [code], <https://github.com/apietak/psit> (last access: 9 April 2026), 2025b.
- Pietak, A.: Data for “psit 1.0: a system to compress Lagrangian flows”, Zenodo [data set], <https://doi.org/10.5281/zenodo.19887527>, 2026.
- Poppick, A., Nardi, J., Feldman, N., Baker, A. H., Pinard, A., and Hammerling, D. M.: A statistical analysis of lossily compressed climate model data, *Comput. Geosci.*, 145, 104599, <https://doi.org/10.1016/j.cageo.2020.104599>, 2020.
- Quach, M., Pang, J., Tian, D., Valenzise, G., and Dufaux, F.: Survey on Deep Learning-Based Point Cloud Compression, *Frontiers in Signal Processing*, 2, <https://doi.org/10.3389/frsip.2022.846972>, 2022.
- Ren, C., Liang, X., and Guo, H.: A Prediction-Traversal Approach for Compressing Scientific Data on Unstructured Meshes with Bounded Error, *Comput. Graph. Forum*, 43, e15097, <https://doi.org/10.1111/cgf.15097>, 2024.
- Rodwell, M., Forbes, R., and Wernli, H.: Why warm conveyor belts matter in NWP, ECMWF newsletter, <https://doi.org/10.21957/mr20vg>, 2018.
- Schemm, S., Wernli, H., and Papritz, L.: Warm conveyor belts in idealized moist baroclinic wave simulations, *J. Atmos. Sci.*, 70, 627–652, <https://doi.org/10.1175/JAS-D-12-0147.1>, 2013.
- Schielicke, L. and Pfahl, S.: European heatwaves in present and future climate simulations: a Lagrangian analysis, *Weather Clim. Dynam.*, 3, 1439–1459, <https://doi.org/10.5194/wcd-3-1439-2022>, 2022.
- Škerlak, B., Sprenger, M., and Wernli, H.: A global climatology of stratosphere–troposphere exchange using the ERA-Interim data set from 1979 to 2011, *Atmos. Chem. Phys.*, 14, 913–937, <https://doi.org/10.5194/acp-14-913-2014>, 2014.
- Skodras, A., Christopoulos, C., and Ebrahimi, T.: The JPEG 2000 still image compression standard, *IEEE Signal Proc. Mag.*, 18, 36–58, <https://doi.org/10.1109/79.952804>, 2001.
- Smith, A. R.: Color gamut transform pairs, *SIGGRAPH Comput. Graph.*, 12, 12–19, <https://doi.org/10.1145/965139.807361>, 1978.
- Sprenger, M. and Wernli, H.: The LAGRANTO Lagrangian analysis tool – version 2.0, *Geosci. Model Dev.*, 8, 2569–2586, <https://doi.org/10.5194/gmd-8-2569-2015>, 2015.
- Sprenger, M., Fragkoulidis, G., Binder, H., Croci-Maspoli, M., Graf, P., Grams, C. M., Knippertz, P., Madonna, E., Schemm, S., Škerlak, B., and Wernli, H.: Global Climatologies of Eulerian and Lagrangian Flow Features based on ERA-Interim, *B. Am. Meteorol. Soc.*, 98, 1739–1748, <https://doi.org/10.1175/BAMS-D-15-00299.1>, 2017.
- Stoffels, R., Benedict, I., Papritz, L., Selten, F., and Weijenborg, C.: Precipitation, moisture sources and transport pathways associated with summertime North Atlantic deep cyclones, *Weather Clim. Dynam.*, 6, 1743–1768, <https://doi.org/10.5194/wcd-6-1743-2025>, 2025.
- Stohl, A., Seibert, P., Wotawa, G., Arnold, D., Burkhart, J. F., Eckhardt, S., Tapia, C., Vargas, A., and Yasunari, T. J.: Xenon-133 and caesium-137 releases into the atmosphere from the Fukushima Dai-ichi nuclear power plant: determination of the source term, atmospheric dispersion, and deposition, *Atmos. Chem. Phys.*, 12, 2313–2343, <https://doi.org/10.5194/acp-12-2313-2012>, 2012.
- Sullivan, G. J., Ohm, J.-R., Han, W.-J., and Wiegand, T.: Overview of the High Efficiency Video Coding (HEVC)

- Standard, IEEE T. Circ. Syst. Vid., 22, 1649–1668, <https://doi.org/10.1109/TCSVT.2012.2221191>, 2012.
- Tintó Prims, O., Redl, R., Rautenhaus, M., Selz, T., Matsunobu, T., Modali, K. R., and Craig, G.: The effect of lossy compression of numerical weather prediction data on data analysis: a case study using enstools-compression 2023.11, *Geosci. Model Dev.*, 17, 8909–8925, <https://doi.org/10.5194/gmd-17-8909-2024>, 2024.
- Unidata: Network Common Data Form (netCDF) version 4.9.2, UniData [code], <https://doi.org/10.5065/D6H70CW6>, 2024.
- Wallace, G. K.: The JPEG still picture compression standard, *Commun. ACM*, 34, 30–44, <https://doi.org/10.1145/103085.103089>, 1991.
- Wendisch, M., Crewell, S., Ehrlich, A., Herber, A., Kirbus, B., Lüpkes, C., Mech, M., Abel, S. J., Akansu, E. F., Ament, F., Aubry, C., Becker, S., Borrmann, S., Bozem, H., Brückner, M., Clemen, H.-C., Dahlke, S., Dekoutsidis, G., Delanoë, J., De La Torre Castro, E., Dorff, H., Dupuy, R., Eppers, O., Ewald, F., George, G., Gorodetskaya, I. V., Grawe, S., Groß, S., Hartmann, J., Henning, S., Hirsch, L., Jäkel, E., Joppe, P., Jourdan, O., Jurányi, Z., Karalis, M., Kellermann, M., Klingebiel, M., Lonardi, M., Lucke, J., Luebke, A. E., Maahn, M., Mahernndl, N., Maturilli, M., Mayer, B., Mayer, J., Mertes, S., Michaelis, J., Michalkov, M., Mioche, G., Moser, M., Müller, H., Neggens, R., Ori, D., Paul, D., Paulus, F. M., Pilz, C., Pithan, F., Pöhlker, M., Pörtge, V., Ringel, M., Risse, N., Roberts, G. C., Rosenburg, S., Röttenbacher, J., Rückert, J., Schäfer, M., Schaefer, J., Schemann, V., Schirmacher, I., Schmidt, J., Schmidt, S., Schneider, J., Schnitt, S., Schwarz, A., Siebert, H., Sodemann, H., Sperzel, T., Spreen, G., Stevens, B., Stratmann, F., Svensson, G., Tatzelt, C., Tuch, T., Vihma, T., Voigt, C., Volkmer, L., Walbröl, A., Weber, A., Wehner, B., Wetzel, B., Wirth, M., and Zinner, T.: Overview: quasi-Lagrangian observations of Arctic air mass transformations – introduction and initial results of the HALO-(AC)³ aircraft campaign, *Atmos. Chem. Phys.*, 24, 8865–8892, <https://doi.org/10.5194/acp-24-8865-2024>, 2024.
- Wernli, H. and Davies, H. C.: A Lagrangian-based analysis of extratropical cyclones. I: The method and some applications, *Q. J. Roy. Meteor. Soc.*, 123, 467–489, <https://doi.org/10.1002/qj.49712353811>, 1997a.
- Wernli, H. and Davies, H. C.: A Lagrangian-based analysis of extratropical cyclones. I: The method and some applications, *Q. J. Roy. Meteor. Soc.*, 123, 467–489, <https://doi.org/10.1002/qj.49712353811>, 1997b.
- Woodring, J., Mniszewski, S., Brislawn, C., DeMarle, D., and Ahrens, J.: Revisiting wavelet compression for large-scale climate data using JPEG 2000 and ensuring data precision, in: 2011 IEEE Symposium on Large Data Analysis and Visualization, 31–38, <https://doi.org/10.1109/LDAV.2011.6092314>, 2011.
- Wotawa, G.: The Fukushima Disaster Calls for a Global Open Data and Information Policy, *GAIA – Ecological Perspectives for Science and Society*, 20, 91–94, <https://doi.org/10.14512/gaia.20.2.5>, 2011.
- Wu, X., Di, S., Ren, C., Jiao, P., Xia, M., Wang, C., Guo, H., Liang, X., and Cappello, F.: Enabling Efficient Error-Controlled Lossy Compression for Unstructured Scientific Data, in: 2025 IEEE International Parallel and Distributed Processing Symposium (IPDPS), 370–382, <https://doi.org/10.1109/IPDPS64566.2025.00040>, 2025.
- Yasunari, T. J., Stohl, A., Hayano, R. S., Burkhart, J. F., Eckhardt, S., and Yasunari, T.: Cesium-137 deposition and contamination of Japanese soils due to the Fukushima nuclear accident, *P. Natl. Acad. Sci. USA*, 108, 19530–19534, <https://doi.org/10.1073/pnas.1112058108>, 2011.
- Zhao, K., Di, S., Dmitriev, M., Tonellot, T.-L. D., Chen, Z., and Cappello, F.: Optimizing Error-Bounded Lossy Compression for Scientific Data by Dynamic Spline Interpolation, in: 2021 IEEE 37th International Conference on Data Engineering (ICDE), 1643–1654, <https://doi.org/10.1109/ICDE51399.2021.00145>, 2021.