Response to Dr. Underwood

Alexander Pietak, Langwen Huang, Luigi Fusco, Michael Sprenger, Sebastian Schemm, and Torsten Hoefler

Dear Dr. Underwood. Thank you for taking the time to provide a review of our submitted manuscript, submitted on 14 Mar 2025. We are pleased your review is supportive of our work and we are of course happy to response to points raised and revise the manuscript accordingly.

The primary improvements we have implemented is the addition of SZ3 as a second baseline comparison to the experiments, along with the inclusion of a rate distortion plot. Additionally, please find below more detailed responses to all your comments.

Comment 1): On line 336, the authors choose a "compression factor of 30". Why was this value chosen? Additionally, the authors should consider multiple values of CRs for this analysis.

Response: Thank you for your suggestion. We have made significant revisions to this experiment, during which we now compare the error distribution of psit with that of ZFP and SZ3 for two different compression ratios of approximately 2.5 and 15. By providing these two compression ratios, we offer a more comprehensive overview of the performance. You may find updated version of the plots in Fig. 1.

Comment 2): This paper misses key references to work on the compression of unstructured grid data which would be more appropriate comparisons for this work

Response: Thank you for the suggestion; we have added a new paragraph in the introduction section where we discuss these papers. Not a large focus is given on these algorithms as they are not inherently designed to compress trajectory data.

Comment 3): The paper should reference the notion of rate distortion in the discussion of Figure 12, and include the PSNR to be more consistent with the standard presentation of rate distortion results, as done in the ZFP and SZ line of papers

Response: Thank you for bringing our attention to this. We have revisited this experiment and have added a PSNR plot. You may find an excerpt of this new analysis in the form of the PSNR plot of the central angle error in Fig. 2.

*Comment 4): Some statements in the paper are not insightful, such as "In general, the L-infinity error is larger than the L1 and RMSE errors." By definition, it is always true that $L_\infty \geq L_1$ because the max upper bounds the average.*Some statements in the paper are not insightful, such as "In general, the L-infinity error is larger than the L1 and RMSE errors." By definition, it is always true that $L_\infty \geq L_1$ because the max upper bounds the average.

Response: We have removed the statement and rewritten parts of the analysis for several experiments.

Comment 5): What justifies the author's "[exception that] the error distribution ... [follows] a Gaussian". This actually depends on the design of the compressor and the error bound chosen. See "Error Distributions of Lossy Floating-Point Compressors" by Lindstrom 2017.

Response: We have reworked this experiment. Our intent is to emphasize that most errors are concentrated at 0.

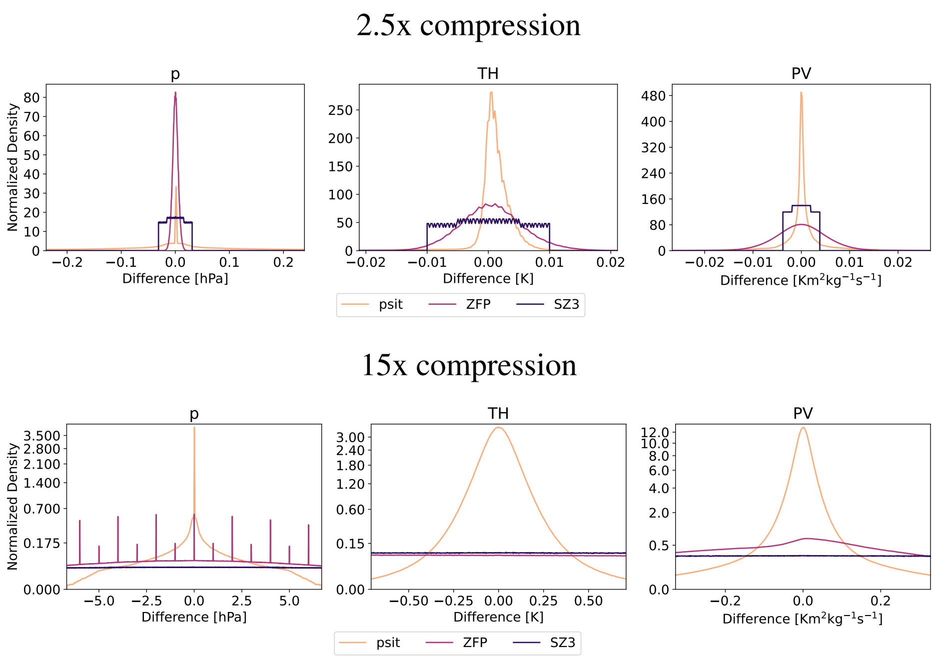

Figure 1. Error distribution based on values for the `tra_20200101_00_permuted` dataset. We compress with compression ratios of 2.5, and 15. The area of the densities have been normalized and the x axis is cropped to 5 standard deviations of the psit distribution for all cases, except pressure with 2.5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

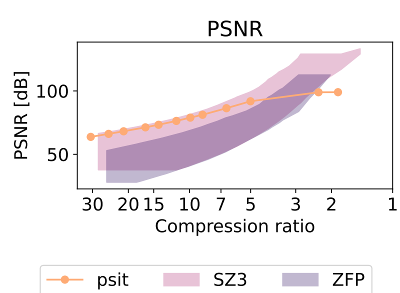

Figure 2. PSNR against compression ratio for the central angle error of the `tra_20200101_00` and the `tra_20200101_00_permuted` trajectory files. The shaded regions for ZFP and SZ3 corresponds to the range in compression performance between the different files.

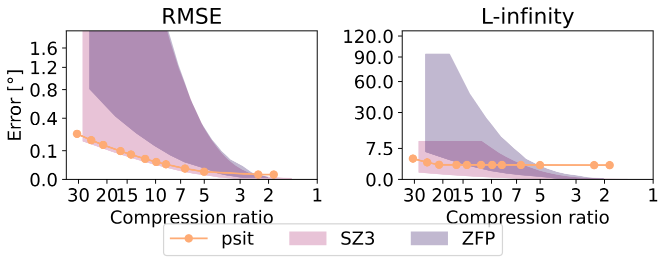

Figure 3. Excerpt of the comparison between psit, ZFP, and SZ3 for the `tra_20200101_00` and `tra_20200101_00_permuted` files. The RMSE and L-infinity error is compared to the achieved compression ratio for the central angle error (lon-lat). In the plots the shaded area for ZFP and SZ3 corresponds to the range in compression performance between the two files. Note that psit performs the same for both, therefore, only one line is plotted.

Comment 6): The limitations of bit rounding-based compressors are well documented in prior work for both SZ and ZFP-based compressors, including papers cited here. It is unclear why this was used as a comparison in section 3.3 of the paper.

Response: Thank you for highlighting this issue. As bitrounding is not the major focus of this experiment and distracted from the comparison, we have taken it out of the manuscript.

Comment 7): It is not clear why the authors compared ZFP and JPEG and not SZ3 when they use ZFP as part of their pipeline in some configurations. It is well known that SZ3 tends to get much higher compression ratios at each quality threshold.

Response: Sorry for the confusion. We have incorporated SZ3 as a second baseline comparison in the experiments, where we are currently only comparing to ZFP. You may find an excerpt of this additional analysis in Fig. 3.

Comment 8): There is extensive work on quantifying appropriate error thresholds by Millian Klöwer and Alison Baker using metrics such as SSIM and dSSIM. Why did the authors not use the metrics in their analysis?

Response: Thank you for your suggestion. We did not include SSIM and dSSIM in the manuscript as they are primarily designed for images (i.e., grid data), but we are focusing on trajectories.

Comment 9): The authors should specify the versions of ZFP and SZ3 used in their work, as newer versions have updated the default algorithm to higher-performing versions. For ZFP, there are versions with much higher parallel performance or support for additional modes with parallel compression.

Response: We have included the versions in the manuscript. For ZFP, it would be pyzfp v0.5.5, which wraps ZFP version 0.5.5, and for SZ3, it is version 3.2.1.

Comment 10): The authors should specify the error bounding type (runtime) and rounding mode (compile time) used with ZFP for their results.

Response: We have added the corresponding information to the manuscript. The bounding type is Fixed-Accuracy. The ZFP version used does not support rounding modes.

Comment 11): E.g., "While this approach is conceptually sound, it quickly, [sic] becomes computationally infeasible, so it found little use in the finished pipeline. However, we still include it here for its theoretical insights." could have been written

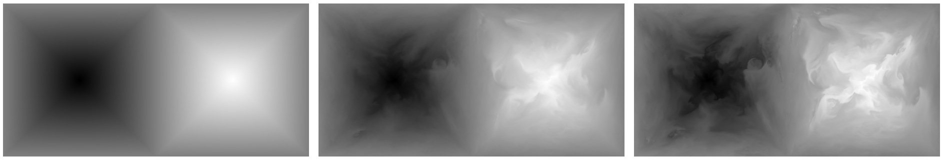

Figure 4. Latitude grids with $550\,\mathrm{hPa}$ starting pressure created with the bipartite mapping method, for detailed information on this refer to section 2.2 of the manuscript. Three different time steps are displayed the leftmost image is at the first time step, the second one is 12 hours later and the rightmost one is 24 hours 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.

"solving an LP is computational infeasible for large images, but included for theoretical analysis."

Response: We have incorporated it into the manuscript.

Comment 12): Additionally, there are many run-on sentences that span 3 or more lines of text.

Response: Thank you for pointing out. We have fixed this.

Comment 13): Lastly, 21 figures (not including subfigures) seems excessive

Response: We have distilled our figures down to the important ones and are now at 13 figures in the main part.

Question 1): On line 330, the authors state, "For longer time ranges, the performance of psit starts to degrade.", Do the authors have an explanation that explains this discrepancy?

Response: The reason is, over time, trajectories start to diverge, which results in the smoothness of the images — dependent on their initial position — being lost. This can be seen in Fig. 4 which corresponds to figure 4 of the manuscript.