

Requirement-Driven Lossy Compression of Weather and Climate Fields

Lagnajeet Roy Lagna

lagnajeetroy@gmail.com

Indian Institute of Technology Hyderabad

Short Report

Keywords:

Posted Date: September 24th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11137790/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

Archives of weather and climate data grow faster than the storage that holds them, and lossy compression is admissible in this domain only when the error it introduces is bounded by an explicit, machine-checkable accuracy requirement. This work addresses five such compression problems on operational atmospheric data: (P1) an absolute error bound on a satellite field in which 69% of the values are missing and must remain missing; (P2) a pointwise relative error bound on precipitation, where exact zeros must be reproduced exactly; (P3) an absolute error bound on the longitudinal derivative of specific humidity rather than on the field itself; (P4) 16 ERA5 pressure-level variables, each carrying its own published safety requirement; and (P5) 226 ERA5 single-level variables, whose requirements span pointwise bounds, mean bounds, data limits, preserved missing values and exact loss lessness. We present a single architecture that solves all five: a requirement analysis that selects a codec family from the *structure* of the bound; a one-parameter search that runs the selected codec progressively past its nominal setting until a reference checker rejects the reconstruction, then bisects; and a shared entropy stage combining tokenisation, bitmap indexing and raw LZMA2. The resulting codecs achieve compression ratios of $\times 47.44$ (P1), $\times 19.07$ (P2) and $\times 83.88$ (P3) with zero violations, and admissible reconstructions for all 242 ERA5 variables with median ratios of $\times 55.68$ (P4) and $\times 50.76$ (P5), against $\times 46.14$ and $\times 22.88$ for the reference safeguards baseline. We further show that admissibility alone is an insufficient acceptance criterion: for 28 of the 226 single-level variables the published mean-error bound exceeds the field's own variability, so the highest-ratio admissible codec reconstructs a constant field, reaching $\times 35,194$ while retaining a single distinct value. We define a fidelity criterion, report it alongside every ratio, and re-derive all affected results under it. A third finding concerns the evaluation code itself: the derivative in P3 is divided by 357.5° rather than the intended 2.5° , leaving that bound roughly $143\times$ looser than its documentation states.

1 Introduction

Numerical weather prediction and climate modelling generate data faster than storage budgets grow. ERA5 [2] occupies petabytes, and kilometer-scale simulations are projected to produce substantially more. Lossless compression of floating-point atmospheric fields typically yields ratios below $\times 2$, while lossy compression yields one to two orders of magnitude more. The obstacle to adopting lossy compression operationally is not the ratio but trust: a reconstruction is usable only if the error it carries is bounded in a form that domain scientists can state, inspect and verify.

Recent work has moved that verification into software. Error-bounded compressors such as ZFP [4], SZ3 [5] and SPERR [6] guarantee a user-specified bound; correction layers [10] enforce a declared

Table 1

The five compression problems addressed. All fields are gridded with atmospheric data; n is the number of variables treated independently.

Field(s)	Grid	Accuracy requirement
P1satellite water-vapour path P2precipitation	$28 \times 320 \times 720$, f32, 69% NaN	$ e \leq 1 \text{ kg m}^{-2}$; NaN preserved
P3specific humidity	$2 \times 721 \times 1440$, f64, 2.7% zeros	$ e \leq 1\% x $ (zeros exact)
	$2 \times 721 \times 1440$, f64	$ e \leq 10^{-6}$ on $\partial x / \partial \lambda$
P4ERA5 pressure level, $n = 16$	$4 \times 721 \times 1440$, f64	per-variable requirement tree
P5ERA5 single level, $n = 226$	721×1440 , f64	per-variable requirement tree

constraint on any codec by storing explicit repairs; and community efforts publish machine-readable accuracy requirements per variable [11] together with a reference implementation that checks them [12]. In this setting the engineering problem changes shape. The compressor is no longer asked to minimize distortion at a given rate; it is asked to maximize the compression ratio subject to a constraint written by somebody else, and to demonstrate admissibility against a checker it does not control.

This manuscript reports a system built for that setting and the five problems it solves, listed in Table 1. The problems were selected because they exercise structurally different constraints: a pointwise bound combined with missing data (P1), a relative bound with exact zeros (P2), a bound on a *derived quantity* rather than on the field (P3), and two large families of per-variable requirements over operational reanalysis data (P4, P5). Our contributions are:

1. **An architecture** that selects a codec family from the structure of the requirement rather than from the data alone and pairs every family with a common entropy stage (Section 2.3).
2. **A checker-in-the-loop parameter search** that exploits the gap between a codec’s worst-case guarantee and its typical error, which is often an order of magnitude (Section 2.4).
3. **A fidelity criterion** that rejects admissible-but-degenerate reconstructions, with evidence that these arise in practice for 28 of 226 variables (Section 2.5).
4. **Quantitative solutions to all five problems** (Section 3), each stored as a codec configuration that is rebuilt from that configuration alone and re-scored automatically.
5. **Two observations about requirement-driven evaluation**: a derivative constraint that is $\approx 143\times$ looser than documented and mean-error bounds that admit erasure of the field (Sections 3.3 and 3.5).

2 Datasets and Methodology

2.1 Datasets

P1: satellite water-vapour path. Six-hourly HOAPS [14] water-vapour path over one week, $28 \times 320 \times 720$ in float32. The field is a satellite swath product: 69% of grid points carry no observation and are stored as NaN, which must remain NaN after reconstruction, while no valid value may become NaN. Valid values lie in $[1.72, 76.69]$ and are quantified to 0.01 in the source data.

P2: precipitation. Two timesteps of total precipitation from an hplp experiment, $2 \times 721 \times 1440$ in float64, non-negative, with 2.7% exact zeros and a strongly skewed distribution spanning twenty orders of magnitude (8.9×10^{-20} to 0.62).

P3: specific humidity. Two timesteps of specific humidity from a NextGEMS ICON simulation [15], $2 \times 721 \times 1440$ in float64, with values in $[2.3 \times 10^{-5}, 2.2 \times 10^{-2}]$.

P4 and P5: ERA5. One timestep of ERA5 [2]: 16 pressure-level variables on four pressure levels ($4 \times 721 \times 1440$, float64) and 262 single-level variables (721×1440 , float64). Of the single-level variables, 226 carry a published requirement and are therefore evaluable; the 36 vertical-integral (vi*) variables have none and are excluded. Fields are streamed from GRIB archives through kerchunk references and cached locally.

2.2 Accuracy requirements

For P1–P3 the requirement is fixed per problem and evaluated by reference code that counts violations over all grid points; a solution is admissible only at zero violations. For P4 and P5 each variable maps to a requirement tree [11] whose leaves are bounds of one of seven types (maximum pointwise absolute, mean absolute, maximum pointwise relative, mean relative, data limits, pre-served missing value, lossless) combined by any and all nodes; the reference checker [12] evaluates the tree against an (original, reconstruction) pair. Two properties of these definitions matter below. A mean bound constrains the average absolute error over finite points, so it tolerates rare large errors but carries no scale relative to the data. A mean *relative* bound constrains total absolute error against total magnitude and requires exact zeros to be reproduced exactly.

2.3 Compression architecture

Figure 1 shows the pipeline. Four codec families cover every requirement encountered across P1–P5.

1. **Zero-anchored uniform grid.** For an absolute bound ϵ , values are snapped to a grid of step 2ϵ anchored at zero and stored in the narrowest integer type that fits. Anchoring at zero is deliberate: fields with many exact zeros (precipitation, snowfall, runoff) reproduce those zeros exactly, which preserves a physically meaningful value and contributes nothing to a mean bound. For mean bounds the step is scaled by the looseness parameter p , because the mean error of uniform rounding lies far below its worst case.
2. **Log-ratio grid.** A pointwise relative bound ϵ is an absolute bound $\log_2(1 + \epsilon)$ on $\log_2 |x|$. A pointwise-ratio meta-codec [9, 13] performs that transform, encodes signs separately and maps exact zeros to

a reserved flag; inside it we place a uniform grid of step just below $2 \log_2(1 + \epsilon)$.

3. **Transform codec, optionally repaired.** SPERR and SZ3 model smooth fields far better than a uniform grid. Their worst-case error equals their nominal setting, but their typical error is much smaller, so under mean bounds they can be run at $p \gg 1$ times the nominal setting and remain admissible. Under hard pointwise bounds they are wrapped in a correction layer [10] that repairs the points exceeding the bound, which is profitable exactly when few points need repair.

Figure 1

The pipeline. A requirement tree is decomposed into leaf bounds; each leaf proposes one or more codec families controlled by a single looseness parameter p . Every family shares the entropy stage, and every candidate is admitted only by the reference checker, with the fidelity criterion applied afterwards.

1. **Correction-layer-only.** A constant (zero) prediction wrapped in the requirement's own safeguards. This is the reference baseline, and it remains the best choice for a small number of variables whose fields are dominated by a constant.

Entropy stage. All families terminate in a common stage: tokenisation of distinct values, a bitmap index that lifts the most frequent token out of the literal stream, and raw LZMA2 with tuned literal-context and position bits. Missing values are handled either by a mask meta-codec storing a packed NaN bitmap or, more effectively, by rounding in floating point so that NaN survives the transform and becomes one more token (Section 3.1).

2.4 Parameter search

Each family exposes one looseness parameter p , a multiple of the requirement's nominal setting. The search starts at $p = 1$, doubles while the checker accepts the result, and on first rejection bisects three times geometrically between the last accepted and the first rejected value. Every evaluated setting is logged with its ratio, admissibility and full configuration, so the selection per variable is a maximum over all sweeps ever run and cannot regress.

Two properties complicate the search. Admissibility is monotone in p , but the ratio is not: for repaired families the ratio rises, peaks and falls again as correction costs overtake the coarser quantization, so for those families we take the maximum over evaluated settings rather than the loosest admissible one. Moreover, repaired candidates are admissible by construction at every p , so the doubling loop never terminates by rejection, and p must be bound explicitly.

2.5 Fidelity criterion

A mean-error bound ϵ is satisfied whenever the average absolute error is below ϵ . If ϵ exceeds the field's own mean magnitude, the constant field satisfies it, and a codec that discards the data entirely is admissible. We therefore record, for every codec, the reconstruction error relative to the field's variability and the number of distinct values retained, and require

RMSE $(x, \hat{x}) \leq 0.30 \sigma(x)$ and $|\{\hat{x}_i\}| > 3, (1)$

in addition to admissibility. The distinct-value guard is deliberately weak: a coarse quantiser retaining only a few dozen levels is legitimate provided its error is small relative to the field, and the guard exists only to catch collapse. Where (1) fails, the highest-ratio logged candidate satisfying both tests is selected instead.

Table 2

Solutions to P1–P3. All have zero violations under the reference evaluation code. Throughput is measured relative to the original size.

PipelineRatio		Size (MB → KB)	Comp.Decomp. (GB/s) (GB/s)	
P1Round → Tokenize → BitmapIndex → LZMA2	47.44	25.8 → 531	0.0227	0.1493
P2LogRatio(grid + BitmapIndex) → LZMA2	19.07	16.6 → 851	0.0124	0.2432
P3Repair(stencil QoI, SPERR q) → LZMA	83.88	16.6 → 193	0.0015	0.1091

2.6 Evaluation protocol

Compression ratio is the ratio of raw to encoded bytes for one encode of the test field, following the reference implementation. A result is reported only at zero violations (P1–P3) or when the reference checker returns true (P4, P5). Throughput is the original size divided by the median wall time of three runs after a warm-up. Every reported codec is serialized to a JSON configuration, rebuilt from that configuration alone and re-scored; this verification runs in continuous integration against checksum-verified input data. All measurements were taken on one workstation (24 logical cores, Python 3.14) using WebAssembly builds of the compressors, so throughput values are comparable within this manuscript but are not statements about native implementations.

3 Results

3.1 P1: absolutely bound with missing values

An absolute bound of 1 on data quantized to 0.01 permits a grid of step 2, which leaves about 40 distinct values. Rounding is performed in floating point so that NaN is unaffected; tokenisation then maps values *and* missingness into a single uint8 stream, and a bitmap index removes the dominant token, which is the swath gap. The result, ×47.44 (Table 2), exceeds every floating-point compressor we applied to this field: SZ3 ×30.3, pcodec ×33.4, bit rounding with Zstandard ×24.3, EBCC ×21.2 and SPERR ×20.1 (Fig. 2a). Representing missingness as a token also beats two obvious alternatives: an explicit NaN bitmap reaches ×39.76, and reconstructing NaN through correction-layer repairs reaches ×36.97. The bitmap index alone accounts for a 10% improvement over the same pipeline without it (×42.80).

3.2 P2: pointwise relative bound with exact zeros

Here the transform codecs lose, and Fig. 3 shows why. Precipitation is spatially noisy, so running SPERR or SZ3 beyond the 1% bound inside the log-ratio codec breaks the bound at 7–23% of points already at $p = 1.41$ and repairing them costs more than the looser setting saves. Even at the nominal setting, where no repair is required, SPERR reaches only $\times 14.2$ and SZ3 $\times 15.8$. A uniform grid in log space, which ignores spatial structure entirely, wins at $\times 19.07$. Adding a bitmap index inside the log-ratio codec contributes a further 0.06%; tokenizing the log bins, which was decisive in P1, is counterproductive here ($\times 13.7$) because the field spans roughly 2,000 distinct bins rather than P1's 40 values.

3.3 P3: bound on a derived quantity

P3 constrains a longitudinal finite difference rather than the field itself. A stencil correction layer encodes that expression exactly, so SPERR can be run at a coarse quantisation step while the layer repairs only those points whose *difference* exceeds the bound. This is the inverse of P2: the errors of a wavelet codec are spatially smooth, so differences between nearby points are largely preserved and few repairs are needed. The sweep in Fig. 2(b) shows a clear optimum at $q = 4.25 \times 10^{-4}$, giving $\times 83.88$; pointwise-bounded alternatives reach only $\times 81.2$ (SPERR), $\times 63.4$ (uniform grid) and $\times 61.6$ (SZ3).

We note that the evaluation code for this problem computes the denominator as `np.mod(lon.roll(5) - lon.roll(-5), 360)`, which evaluates to 357.5 at every point, whereas the separation between the two sampled points is 2.5° . The constraint enforced is therefore roughly 143 $\times$ looser than the accompanying documentation states. Our solution targets the constraint as implemented, as any evaluated submission must, and we report the discrepancy rather than depend on it silently.

3.4 P4 and P5: per-variable requirements over ERA5

Table 3 and Fig. 4 summarize all 242 variables, every one of which is admissible and satisfies the fidelity criterion. The highest ratios are $\times 787.7$ for temperature and $\times 135.9$, $\times 132.2$ and $\times 131.9$ for the wind components and geopotential (P4), and $\times 6,530$ for ultraviolet albedo (P5); the lowest are $\times 13.5$ for divergence (P4) and $\times 8.1$ for lake-depth field (P5).

The winning family differs sharply between the two problems. Among pressure-level variables, transform codecs win 9 of 16 and the log-ratio grid 5; among single-level variables the uniform grid wins 170 of 226, with SPERR 23, log-ratio 13, correction-layer-only 11, lossless 5 and SZ3 4. The pattern tracks the requirement type more closely than the field itself: pressure-level fields are three-dimensional, smooth, and carry pointwise bounds, which Favors transform coding, whereas single-level fields are dominated by mean bounds, under which a zero-anchored grid at a large step is difficult to beat, particularly for the many zero-inflated surface fields.

Table 3

Solutions to P4 and P5, after the fidelity criterion is applied. “Baseline” is the reference correction-layer codec. The aggregate ratio is total raw bytes over total compressed bytes across a level’s variables.

Variables admissible	P4: pressure level (16)	P5: single level (226)
	16/16	226/226
Median ratio	55.68	50.76
Median baseline ratio	46.14	22.88
Median gain over baseline	1.35	1.85
Variables $\geq 2\times$ baseline	4	100
Aggregate ratio	44.10	43.00
Median RMSE / $\sigma(\text{field})$	0.006	0.008
Median decompression (GB/s)	0.074	0.506

3.5 Admissibility is not sufficient: the fidelity criterion

For 28 of the 226 single-level variables the requirement’s mean-error bound exceeds the field’s own mean magnitude, and the highest-ratio admissible codec consequently reconstructs a (near) constant field: $\times 35,194$ for csf, es and smlt by returning a single distinct value, and $\times 34,899$ for istl4 with an error $2.3\times$ the field’s standard deviation (Table 4, Fig. 5a). These codecs satisfy the published requirement exactly as written.

Applying criterion (1) costs a substantial amount of headline ratio: the maximum single-level ratio falls from $\times 35,194$ to $\times 6,530$, and the count of variables above $\times 1000$ falls from 27 to 20. We consider this the correct trade, since reconstruction retaining one distinct value is not a compressed field. After the correction, the worst reconstruction error across all 242 variables is 0.30σ and the median is 0.008σ . A related implementation detail is worth recording: translating a data-limit requirement (for example, non-negativity) into safeguards returns an empty safeguard list when that requirement appears alone, so a codec that undershoots the limit is rejected by the checker rather than repaired. Enforcing it instead with a sign-preserving safeguard is possible but expensive on zero-inflated fields: on convective precipitation it costs a factor of 37 in ratio ($\times 141.66$ to $\times 3.83$), because each of the 62% exact zeros then requires a correction.

3.6 Throughput

Figure 5(b) relates ratio to decompression speed. The uniform grid pipelines decompress fastest (up to 1.4 GB/s of original size), because decoding is a table lookup followed by a multiply; transform codecs are slower; and repaired pipelines are slowest, since the correction layer evaluates the constraint at

every point. Compression is consistently slower than decompression, with medians of 0.008 GB/s (P4) and 0.015 GB/s (P5). For the archival use case that motivates this work, where data is written once and read many times, this asymmetry is favorable.

3.7 Comparison with independent solutions

The laboratory that publishes these problems also collects independent solutions. At the time of writing (23 September 2026), published entries for P1–P3 stood at $\times 68.96/\times 53.92/\times 47.44/\times 40.18$, $\times 21.76/\times 19.07/\times 18.07$ and $\times 101.27/\times 83.88/\times 82.12$ respectively, ours being $\times 47.44$, $\times 19.07$ and $\times 83.88$; for P4 and P5, on the two variables where a comparison exists, other entries report $\times 141.87$ and $\times 143.84$ against our $\times 131.89$ and $\times 141.66$. Several of the higher entries report an external tool and a single configuration line (compression level = 5), which cannot be rebuilt from what is published, so neither their admissibility under the same checker nor their reconstruction fidelity can be assessed independently. This is precisely the situation that the configuration-as-artefact convention is intended to prevent.

Table 4: The ten largest corrections made by the fidelity criterion in P5. “Distinct” is the number of distinct values in the reconstruction. Every row on both sides satisfies the variable’s published requirement.

Variable	Ratio	RMSE/ σ	Distinct	Ratio	RMSE/ σ	Distinct
es	35194.58	1.00	1	717.26	0.05	114
csf	35194.58	1.00	1	541.84	0.04	84
smlt	35194.58	1.00	1	5732.17	0.27	8
istl4	34898.82	2.30	1	484.68	0.01	151
lgws	34040.66	1.00	2	177.57	0.06	215
avg_esrwe	33091.32	1.00	2	3208.16	0.30	14
mgws	30536.47	0.96	3	184.27	0.06	198
avg_smr	21351.98	0.88	2	3498.70	0.12	19
tcrw	15129.18	0.78	3	1433.29	0.18	32
ro	14546.27	0.70	4	1171.50	0.22	35

3.8 Limitations

The ERA5 measurements use one-timestep test subsets; ratios over a full month may differ, particularly for transform codecs, which benefit from the additional dimension. Throughput was measured on a single machine with WebAssembly builds. The parameter search is a one-dimensional sweep per family

rather than a joint optimisation over codec and entropy-stage settings, and the full search was run for all 16 pressure-level variables but only 46 of 226 single-level ones, so the remainder use a coarser sweep and are likely improvable. The thresholds in (1) are a judgement rather than a derived criterion; the underlying quantities are reported per variable so that any other threshold can be applied.

4 Conclusion

We have described a single architecture that solves five structurally different compression problems on operational atmospheric data and reported its results under the reference checkers that define admissibility. Selecting the codec family from the structure of the requirement, searching for the looseness parameter with the checker in the loop, and sharing one entropy stage across families yields $\times 47.44$ for an absolute bound on a 69%-missing satellite field, $\times 19.07$ for a pointwise relative bound on precipitation, $\times 83.88$ for a bound on a longitudinal derivative, and admissible reconstructions for all 242 ERA5 variables at median ratios of $\times 55.68$ and $\times 50.76$, roughly 1.35 and 1.85 times the reference baseline.

Three findings generalize beyond these datasets. First, the choice between a transform codec and a uniform quantiser is governed by the interaction of field regularity with the bound type: smooth fields under pointwise bounds favor transform coding, while noisy or zero-inflated fields under mean bounds favor a zero-anchored grid whose exact zeros are free. Second, a correction layer converts any codec into an admissible one, but its cost scales with the fraction of points repaired, so the violated fraction at the intended setting is a reliable predictor of whether repair will pay: it does for a derivative constraint on a smooth field (P3) and does not for a pointwise constraint on a noisy one (P2). Third, and most consequential for benchmark design, admissibility alone is not a sufficient acceptance criterion. When a mean-error bound exceeds a field's own variability, erasing the field is admissible, and we observed this for 28 of 226 variables at ratios up to $\times 35,194$. We therefore recommend that requirement-driven compression results be reported with a fidelity statement alongside the ratio, and that every reported codec be published as a configuration a third party can rebuild and re-score. Both conventions are followed throughout this work.

Declarations

Code and data availability

All code, codec configurations, sweep logs and result tables are available at <https://github.com/Lagnajee/compression-hackathon-solutions>. The datasets are those distributed by the com-pression laboratory [1]; scripts/download_data.sh fetches and checksum-verifies them. The figures in this manuscript are regenerated from the result files by paper/figures.py.

References

1. ESiWACE3/ECMWF. *Online Laboratory for Data Compression in Climate Science and Meteorology*. <https://github.com/climet-eu/compression-lab-notebooks>.
2. H. Hersbach et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.
3. Zarr Developers. *numcodecs*: a collection of compression codecs for buffer compression. <https://github.com/zarr-developers/numcodecs>, version 0.15.1.
4. P. Lindstrom. Fixed-rate compressed floating-point arrays. *IEEE Transactions on Visualization and Computer Graphics*, 20(12), 2014.
5. X. Liang et al. SZ3: a modular framework for composing prediction-based error-bounded lossy compressors. *IEEE Transactions on Big Data*, 2023.
6. S. Li et al. Lossy scientific data compression with SPERR. *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2023.
7. A. Fallin and M. Burtscher. Lessons learned on the path to guaranteeing the error bound in lossy quantizers. *arXiv:2407.15037*, 2024.
8. M. Klöwer et al. Compressing atmospheric data into its real information content. *Nature Computational Science*, 1:713–724, 2021.
9. X. Liang, S. Di, D. Tao, Z. Chen, and F. Cappello. An efficient transformation scheme for lossy data compression with point-wise relative error bound. *IEEE International Conference on Cluster Computing (CLUSTER)*, 179–189, 2018.
10. J. Tyree. *compression-safeguards*. <https://github.com/juntyr/compression-safeguards>, version 1.0.0rc3.
11. *compression-recommendations*: machine-readable safety requirements for atmospheric variables. <https://github.com/juntyr/compression-recommendations>, version 0.1.0a1.
12. *compression-requirement-checks*: reference implementation of the safety requirement checks. <https://github.com/juntyr/compression-recommendations>, version 0.1.0a1.
13. *numcodecs-pw-ratio*: pointwise ratio error bounded meta-codec. <https://pypi.org/project/numcodecs-pw-ratio/>, version 0.2.1.
14. A. Andersson et al. The Hamburg Ocean Atmosphere Parameters and Fluxes from Satellite Data (HOAPS): a climatological atlas of satellite-derived air-sea-interaction parameters over the oceans. *Earth System Science Data*, 2010.
15. nextGEMS: Next Generation Earth Modelling Systems. <https://nextgems-h2020.eu/>.
16. T. Reichelt, J. Tyree, M. Klöwer, P. Dueben, B. N. Lawrence, A. H. Baker, S. Faghih-Naini, T. Hoefler, and P. Stier. ClimateBenchPress (v1.0): a benchmark for lossy compression of climate data. *Geoscientific Model Development*, 19:5933, 2026.

Figures




Figure 1

The pipeline. A requirement tree is decomposed into leaf bounds; each leaf proposes one or more codec families controlled by a single looseness parameter p . Every family shares the entropy stage, and every candidate is admitted only by the reference checker, with the fidelity criterion applied afterwards.





Figure 2

(a) P1: admissible ratios for the pipelines evaluated. Blue bars are variants of the proposed architecture; grey bars are established compressors applied to the same field. (b) P3: ratio against codec setting under the stencil-Qol correction layer; the selected SPERR quantisation step lies at the optimum.

(a) P2: transform codecs stay below the quantiser



(b) P2: why repair is uneconomic





Figure 3

P2. (a) Ratio against looseness for transform codecs inside the log-ratio codec, with and without a relative-error correction layer; circles are admissible, crosses are rejected. The horizontal line is the uniform log grid. (b) The fraction of points violating the bound grows rapidly with looseness, which is what makes repair uneconomic for this field.

(a) P4/P5: distribution of ratios



(b) P4/P5: gain over the baseline




Figure 4

(a) Distribution of compression ratios for P4 and P5. (b) Each variable against the reference baseline; points above the dotted line are at least twice the baseline.




Figure 5

(a) Compression ratio against reconstruction error for P5. Crosses are the admissible-but-degenerate results; arrows show where the fidelity criterion moves them. (b) Ratio against decompression throughput, coloured by winning family; repaired pipelines are the slow ones.