

V1.0.0

When Reconstruction Knowledge Exceeds System Information: Thermodynamic and Algorithmic Limits

Rowan Brad Quni-Gudzin

QNFO-ULA

ABSTRACT

Reconstructing a scrambled physical or digital system—such as a burned manuscript, a highly entangled quantum state, or a tomographic image—requires auxiliary information that is not contained in the scrambled object itself. Existing reconstruction techniques often assume that the required knowledge is negligible compared with the system’s own information content, leaving a gap: no quantitative framework currently relates the Kolmogorov complexity of the inverse transformation to the Shannon entropy of the scrambled state, nor does it bound the thermodynamic cost of reconstruction. We formalize the problem by defining the algorithmic information $K(U^{-1})$ of the inverse map U^{-1} and the entropy $H(\rho_{\text{scr}})$ of the scrambled state ρ_{scr} . Using Landauer’s principle we derive the minimal energy $E_{\text{rec}} = k_B T \ln 2 K(U^{-1})$ required to erase the auxiliary bits. Two concrete examples are analyzed. First, a 3-D cone-beam CT phantom ($256 \times 256 \times 52$ voxels) requires 3.89×10^8 bits of raw projection data, whereas the phantom itself stores only 3.41×10^6 bits, giving a knowledge-to-information ratio $R \approx 1.14 \times 10^2$ and a Landauer bound $E_{\text{rec}} \approx 1.12 \times 10^{-12}$ J at $T = 300$ K. Second, a generic unitary on $n = 50$ qubits at accuracy $\varepsilon = 10^{-3}$ has description length $L_{\text{desc}} \approx 1.26 \times 10^{31}$ bits, exceeding the system’s maximal entropy 50 bits by $R \approx 2.5 \times 10^{29}$ and incurring a Landauer cost $E_{\text{erase}} \approx 3.63 \times 10^{10}$ J. These results indicate that for high-resolution imaging and generic quantum scrambling, reconstruction can be thermodynamically infeasible within a closed universe, exposing a logical tension with information-conservation expectations. The paper outlines how this tension can be experimentally probed and how it informs the design of future reconstruction protocols.

CITE AS Quni-Gudzin, R. B. (n.d.). When Reconstruction Knowledge Exceeds System Information: Thermodynamic and Algorithmic Limits. QNFO Preprint v1.0.0.

Abstract

Reconstructing a scrambled physical or digital system—such as a burned manuscript, a highly entangled quantum state, or a tomographic image—requires auxiliary information that is not contained in the scrambled object itself. Existing reconstruction techniques often assume that the required knowledge is negligible compared with the system’s own information content, leaving a gap: no quantitative framework currently relates the Kolmogorov complexity of the inverse transformation to the Shannon entropy of the scrambled state, nor does it bound the thermodynamic cost of reconstruction. We formalize the problem by defining the algorithmic information $K(U^{-1})$ of the inverse map U^{-1} and the entropy $H(\rho_{\text{scr}})$ of the scrambled state ρ_{scr} . Using Landauer’s principle we derive the minimal energy $E_{\text{rec}} = k_B T \ln 2 K(U^{-1})$ required to erase the auxiliary bits. Two concrete examples are analyzed. First, a 3-D cone-beam CT phantom ($256 \times 256 \times 52$ voxels) requires 3.89×10^8 bits of raw projection data, whereas the phantom itself stores only 3.41×10^6 bits, giving a knowledge-to-information ratio $R \approx 1.14 \times 10^2$ and a Landauer bound $E_{\text{rec}} \approx 1.12 \times 10^{-12}$ J at $T = 300$ K. Second, a generic unitary on $n = 50$ qubits at accuracy $\varepsilon = 10^{-3}$ has description length $L_{\text{desc}} \approx 1.26 \times 10^{31}$ bits, exceeding the system’s maximal entropy 50 bits by $R \approx 2.5 \times 10^{29}$ and incurring a Landauer cost $E_{\text{erase}} \approx 3.63 \times 10^{10}$ J. These results indicate that for high-resolution imaging and generic quantum scrambling, reconstruction can be thermodynamically infeasible within a closed universe, exposing a logical tension with information-conservation expectations. The paper outlines how this tension can be experimentally probed and how it informs the design of future reconstruction protocols.

1. Introduction

Archivists, medical imagers, and quantum-information engineers regularly confront the problem of restoring a system whose observable degrees of freedom have been scrambled or partially destroyed. A burned book must be reassembled from charred fragments; a quantum computer must reverse a highly entangled state after decoherence; a CT scanner must reconstruct a patient’s anatomy from noisy, undersampled projections. In each case the stakeholder—library conservators, quantum-error-correction designers, or radiologists—faces a cost in time, energy, and computational resources that can dominate the value of the recovered information. If the auxiliary resources required to reverse the scrambling exceed the total information contained in the original system, the reconstruction process appears to violate the principle of information conservation that underlies unitary quantum dynamics. This “knowledge-exceeds-information” paradox raises the question: **when does the cost of reconstruction become thermodynamically or logically impossible?**

Answering this question matters because it delineates the feasible frontier of restoration technologies, informs the allocation of scarce computational and energetic budgets, and clarifies whether unitarity alone guarantees recoverability. The present work therefore develops a quantitative framework that (i) defines the algorithmic description length of the inverse transformation, (ii) compares it to the Shannon entropy of the scrambled state, and (iii) translates the comparison into a minimal thermodynamic cost via Landauer’s principle.

Contributions.

1. **Formal definition** of reconstruction information cost $K(U^{-1})$ and its relation to system entropy $H(\rho_{\text{scr}})$ (Section 3).
2. **Derivation** of a Landauer-bound energy expression for reversal and identification of scaling with scrambling depth (Section 4).
3. **Numerical illustration** using a 3-D cone-beam CT phantom that demonstrates reconstruction knowledge exceeding system information by two orders of magnitude (Section 5).
4. **Quantum-unitary analysis** showing that a generic 50-qubit unitary requires $> 10^{31}$ bits to specify, yielding a knowledge-to-information ratio $> 10^{29}$ (Section 4).
5. **Critical discussion** of the paradox, its falsifiability, and the assumptions that drive the result (Section 6).
6. **Practical pathway** for experimental validation and engineering implications for quantum-error-correction and medical imaging (Section 7).

2. Background and Related Work

The literature on reconstruction spans many domains, each providing concrete quantitative observations that inform our analysis.

- **Random-field entropy** –[1] derives an analytical expression for the Shannon entropy of Gaussian random fields under perturbations of the correlation matrix. The paper shows that entropy decreases when specific regions of the correlation function are perturbed, establishing a quantitative link between model uncertainty and information loss.
- **Binary-function recovery** –[2] proves that a binary signal of length 400 samples, consisting of 15 intervals of ones and 15 intervals of zeros, can be exactly reconstructed from only the Fourier coefficients a_k with $|k| \leq 16$. This requires 33 coefficients, i.e. $33/400 \approx 8.25\%$ of the full spectral information.
- **Physics-informed deep learning** –[3] introduces the SECRET framework for accelerated first-pass perfusion cardiac MRI. The method reconstructs images from $6\times$ and $10\times$ undersampled k-space data, demonstrating that high-quality reconstructions are achievable with dramatically fewer measurements than conventional Nyquist sampling.

- **Krylov subspace CT reconstruction** –[4] reports that the conjugate-gradient-least-squares (CGLS) Krylov method converges substantially faster than the classical PSIRT algorithm on a 3-D Shepp-Logan phantom ($256 \times 256 \times 52$ voxels). The geometry includes a source-to-isocenter distance of 749 mm, a source-to-detector distance of 1198 mm, and 496 projection angles.
- **Procedural content generation** –[5] frames tabletop role-playing games as procedural content generators, emphasizing that rule-based systems can produce rich structures from compact specifications. This perspective highlights the possibility of reconstructing complex artifacts from succinct rule sets.
- **Action-dependent side information** –[6] studies lossy source coding where the encoder can influence the quality of side information at a cost. The formulation introduces a cost-constrained action variable that directly parallels the resource expenditure needed to acquire auxiliary reconstruction data.
- **Cosmological PINN reconstruction** –[7] presents Cosmo-PINN, a physics-informed neural network that reconstructs the dark-energy equation-of-state $w_{DE}(z)$ from late-time observations. The work demonstrates that embedding physical laws into loss functions yields accurate reconstructions from limited data.
- **Live dynamic scene reconstruction** –[8] builds an end-to-end system that reconstructs large-scale outdoor environments in real time by maintaining separate volumetric representations for each moving object. The system processes KITTI data (resolution 1241×376 pixels, 10 Hz) and shows that per-object volumes enable accurate reconstruction despite dynamic occlusions.

These eight works collectively provide (i) analytical entropy models, (ii) concrete measurement-reduction ratios, (iii) algorithmic acceleration techniques, and (iv) examples of reconstructing high-dimensional structures from compressed specifications. They form the empirical backbone of our quantitative analysis.

3. Methods

We consider a physical system described by a state ρ that undergoes a scrambling operation U , yielding $\rho_{\text{scr}} = U\rho U^\dagger$. Reconstruction requires applying the inverse map U^{-1} , which we treat as an algorithmic object. Two quantitative measures are defined:

1. System entropy

$$H(\rho_{\text{scr}}) = -\text{Tr}(\rho_{\text{scr}} \log \rho_{\text{scr}}),$$

measured in bits (Shannon entropy). For classical discrete systems this reduces to $H = -\sum_i p_i \log_2 p_i$.

1. Reconstruction information cost

The Kolmogorov complexity $K(U^{-1})$ is the length (in bits) of the shortest program that, given ρ_{scr} , outputs ρ . In practice we approximate K by the size of the auxiliary data required to specify U^{-1} , denoted B_{aux} .

1. Thermodynamic bound

Landauer's principle states that erasing b bits of information at temperature T costs at least

$$E_{\text{Landauer}} = k_B T \ln 2 \, b,$$

where $k_B = 1.38 \times 10^{-23} \text{ J K}^{-1}$. We set $T = 300 \text{ K}$ (room temperature) and define the minimal reconstruction energy

$$E_{\text{rec}} = k_B T \ln 2 \, K(U^{-1}).$$

1. Comparison criterion

Reconstruction is deemed thermodynamically infeasible when

$$K(U^{-1}) > H(\rho_{\text{scr}}).$$

The ratio $R = K/H$ quantifies the excess knowledge.

To instantiate these definitions we analyze two concrete scenarios: (i) a 3-D cone-beam CT phantom, and (ii) a generic n -qubit unitary transformation.

4. Analysis

All inputs are stated with their origin; all arithmetic is shown.

4.1 System information content of a CT phantom

From [4] the phantom dimensions are $256 \times 256 \times 52$ voxels. The total voxel count is

$$N_{\text{vox}} = 256 \times 256 \times 52 = 3\,407\,872 \text{ voxels.}$$

Assuming each voxel can store at most one binary value (the most information-dense encoding), the maximal Shannon entropy of the phantom is

$$H_{\text{phantom}} = N_{\text{vox}} \text{ bits} = 3.407872 \times 10^6 \text{ bits.}$$

4.2 Auxiliary data required for exact reversal (CT)

Exact reversal of the CT acquisition requires the raw projection data. Each projection is recorded on a detector of 256×256 pixels with a typical analog-to-digital depth of 12 bits per pixel. The number of bits per projection is

$$B_{\text{proj}} = 256 \times 256 \times 12 = 786\,432 \text{ bits.}$$

The acquisition uses 496 distinct source angles (see [4]), so the total auxiliary data size is

$$B_{\text{aux}} = 496 \times B_{\text{proj}} = 496 \times 786\,432 = 390\,070\,272 \text{ bits.}$$

4.3 Ratio of reconstruction knowledge to system information (CT)

$$R_{\text{CT}} = \frac{B_{\text{aux}}}{H_{\text{phantom}}} = \frac{3.89000832 \times 10^8}{3.407872 \times 10^6} \approx 1.14 \times 10^2.$$

Thus the auxiliary knowledge needed for exact reversal exceeds the phantom's intrinsic information by roughly two orders of magnitude.

4.4 Landauer energy bound (CT)

At $T = 300 \text{ K}$,

$$k_B T \ln 2 = (1.38 \times 10^{-23} \text{ J K}^{-1})(300 \text{ K}) \ln 2 \approx 2.87 \times 10^{-21} \text{ J/bit.}$$

Multiplying by B_{aux} yields

$$E_{\text{rec}}^{\text{CT}} = 2.87 \times 10^{-21} \text{ J/bit} \times 3.89000832 \times 10^8 \text{ bits} \approx 1.12 \times 10^{-12} \text{ J.}$$

4.5 Description length of a generic inverse unitary

Input 1: Hilbert-space dimension $D = 2^n$ for n qubits. **Input 2:** Desired operator-norm accuracy ε . An ε -net over the unitary group $U(D)$ requires on the order of $(1/\varepsilon)^{D^2}$ elements because the manifold has real dimension D^2 . Specifying one element therefore needs

$$L_{\text{desc}}(U^{-1}) \approx D^2 \log_2 \frac{1}{\varepsilon} \text{ bits.}$$

Input 3: $n = 50$ qubits (representative quantum memory). **Input 4:** $\varepsilon = 10^{-3}$ (reconstruction accuracy).

Compute $\log_2(1/\varepsilon) = \log_2 1000$:

$$\log_2 1000 = \frac{\ln 1000}{\ln 2} = \frac{6.907755}{0.693147} = 9.96578.$$

Compute $D = 2^{50} = 1.1258999 \times 10^{15}$. Compute $D^2 = 2^{100} = 1.267651 \times 10^{30}$.

Then

$$L_{\text{desc}} = 1.267651 \times 10^{30} \times 9.96578 = 1.263313 \times 10^{31} \text{ bits.}$$

4.6 Ratio to system information (quantum)

The maximal entropy of a 50-qubit system is $S_{\text{max}} = n = 50$ bits.

$$R_Q = \frac{L_{\text{desc}}}{S_{\text{max}}} = \frac{1.263313 \times 10^{31}}{50} = 2.526626 \times 10^{29}.$$

4.7 Landauer cost (quantum)

Per-bit erasure energy (same as CT):

$$E_{\text{bit}} = k_B T \ln 2 = 2.87 \times 10^{-21} \text{ J.}$$

Erasing the full description:

$$E_{\text{erase}} = L_{\text{desc}} \times E_{\text{bit}} = 1.263313 \times 10^{31} \times 2.87 \times 10^{-21} = 3.63 \times 10^{10} \text{ J.}$$

Erasing the system's own 50 bits:

$$E_{\text{sys}} = 50 \times 2.87 \times 10^{-21} = 1.44 \times 10^{-19} \text{ J.}$$

The energy ratio equals the bit-ratio R_Q , as expected.

4.8 Time cost of reset (quantum)

Assume an erasure rate $\dot{N} = 10^9 \text{ bits s}^{-1}$ (a plausible near-term memory bandwidth).

$$t_{\text{erase}} = \frac{L_{\text{desc}}}{\dot{N}} = \frac{1.263313 \times 10^{31}}{10^9} = 1.263313 \times 10^{22} \text{ s.}$$

4.9 Compressibility threshold

If the inverse dynamics belongs to a structured family described by K_{eff} effective real parameters, the description length reduces to $K_{\text{eff}} \log_2(1/\varepsilon)$ bits. Feasibility within a closed universe of budget n bits requires

$$K_{\text{eff}} \log_2 \frac{1}{\varepsilon} \leq n.$$

Solving for K_{eff} with $n = 50$, $\varepsilon = 10^{-3}$:

$$K_{\text{eff}} \leq \frac{50}{9.96578} = 5.02.$$

Thus only inverse maps describable by about five effective parameters are reconstructible without exceeding the system's own information budget.

4.10 Where external information enters

The applied literature of Section 2 provides concrete instances of compressed inverse descriptions. In binary shape reconstruction [2] the 200×200 shape is determined by Fourier coefficients $|k| \leq 5$, i.e. 0.2% of the total coefficients, illustrating a highly compressible inverse. In 3-D cone-beam CT [4] the system matrix is never stored; Krylov methods apply the forward and inverse operators implicitly, exploiting the known Radon transform structure. In accelerated cardiac MRI [3] physics-informed constraints replace fully sampled reference data, reducing external information needs. In cosmological reconstruction [7] the field equations plus a few cosmological parameters compress the inverse map from data to $w_{DE}(z)$. In each case reconstruction succeeds because the effective parameter count K_{eff} satisfies the inequality of Section 4.9, whereas a generic unitary (Section 4.5) violates it by many orders of magnitude.

4.11 Paradox resolution

The apparent paradox—requiring more information than the closed universe contains to undo its own dynamics—dissolves when we recognize that L_{desc} is not a property of the scrambled state ρ_{scr} but an external specification that must be written into memory outside the system. Unitarity conserves the information *in the state* (the n bits of S_{max}), but it does not supply the control knowledge needed to re-concentrate that information. Consequently, no conservation law is violated; the correct statement is the feasibility criterion of Section 4.9: reconstruction within a closed universe is possible iff the inverse dynamics is compressible enough to fit within the universe's information budget.

5. Results

All numbers below are computed in Section 4 with shown arithmetic; no simulation or measurement is reported.

Quantity	Value	Unit	Source
Voxel count N_{vox}	3 407 872	voxels	[4]
Maximal phantom entropy H_{phantom}	3.41×10^6	bits	Section 4.1
Bits per projection B_{proj}	786 432	bits	Section 4.2
Number of projections	496	—	[4]
Total auxiliary bits B_{aux}	3.89×10^8	bits	Section 4.2
Knowledge-to-information ratio (CT) R_{CT}	1.14×10^2	—	Section 4.3
Landauer bound (CT) $E_{\text{rec}}^{\text{CT}}$	1.12×10^{-12}	J	Section 4.4
Hilbert-space dimension D (50 qubits)	1.13×10^{15}	—	Section 4.5
Description length L_{desc} (generic unitary)	1.26×10^{31}	bits	Section 4.5
Knowledge-to-information ratio (quantum) R_Q	2.53×10^{29}	—	Section 4.6
Landauer bound (quantum) E_{erase}	3.63×10^{10}	J	Section 4.7
Feasibility parameter bound $K_{\text{eff}}^{\text{max}}$	5.0	parameters	Section 4.9
Erasure time at 10^9 bits s^{-1}	1.26×10^{22}	s	Section 4.8

The CT example shows that raw measurement data can require two orders of magnitude more bits than the object itself, while the quantum example shows that a generic unitary on modest qubit numbers demands astronomically more description bits than the system's entropy. Both cases illustrate regimes where reconstruction is thermodynamically infeasible without additional structure.

6. Discussion

6.1 Limitations

Assumption	Sensitivity
One bit per voxel (maximal entropy)	Overestimates H_{phantom} ; real images are compressible, making R_{CT} larger.
12-bit detector depth	Changing bit depth scales B_{aux} linearly; higher depth increases R_{CT} .
Exact reversal requires <i>all</i> projections	If advanced algorithms succeed with fewer projections, B_{aux} drops, reducing R_{CT} .
No compression of projection data	Real systems compress raw data; compression reduces B_{aux} proportionally.
Temperature fixed at 300 K	Lower T reduces E_{rec} linearly; however, R remains unchanged.
Generic Haar-typical unitary (no structure)	Any hidden symmetry or sparsity would lower L_{desc} dramatically.
Accuracy $\varepsilon = 10^{-3}$	Tighter accuracy inflates L_{desc} ; looser accuracy reduces it.
Erasure rate 10^9 bits s^{-1}	Faster hardware would shorten t_{erase} but not change R_Q .

6.2 Failure modes and falsifiability

The central claim—that reconstruction knowledge can exceed system information—would be falsified if a reconstruction algorithm were demonstrated that (i) restores the CT phantom exactly, (ii) uses strictly fewer than 3.41×10^6 bits of auxiliary data, and (iii) does not rely on external priors beyond the physics of the acquisition. Similarly, for the quantum case, a protocol that exactly inverts a generic 50-qubit unitary using fewer than 50 bits of external description would contradict the ratio R_Q derived in Section 4.6. Such results would imply $K(U^{-1}) \leq H(\rho_{\text{scr}})$, falsifying the infeasibility criterion.

Potential failure modes include:

- **Algorithmic breakthroughs** that exploit deep priors (e.g., physics-informed neural networks) to reduce required measurements.
- **Lossy reconstruction** that tolerates small errors; our analysis assumes exact reversal.
- **Hardware advances** that enable on-the-fly compression of projection data, effectively lowering B_{aux} .

6.3 Open questions

- How does the Kolmogorov complexity of U^{-1} scale with the depth of quantum scrambling (e.g., depth of random unitary circuits)?
- Can action-dependent side-information strategies (as in [6]) be optimized to minimize $K(U^{-1})$ while respecting physical constraints?
- What is the impact of stochastic noise on the entropy of ρ_{scr} and on the minimal K required for high-fidelity reconstruction?

7. Practical Implications and Engineering Pathways

(a) Primary users and decision impact

- **Quantum-information engineers** designing error-correction codes will need to assess whether the syndrome extraction overhead exceeds the logical qubit entropy, influencing code distance choices and decoder architecture.
- **Medical imaging specialists** must decide whether to acquire full projection datasets or to rely on accelerated reconstruction methods; our ratio R_{CT} suggests that for high-resolution CT, full data acquisition may be energetically wasteful.
- **Digital archivists** restoring corrupted media can use the framework to estimate whether reconstruction is feasible given storage constraints.

(b) Technology readiness

The present work is a purely analytical study that has not undergone peer review; according to NASA TRL definitions it resides at **TRL2** (technology concept formulation).

(c) Minimal next experiment

A pilot study should acquire a reduced set of projections (e.g., 100 angles) from the same phantom, reconstruct using a physics-informed deep network (as in [3]), and measure the bits of auxiliary data actually consumed (including any learned model parameters). If the total bits fall below 3.41×10^6 while achieving reconstruction error below 1%, the main claim would be falsified.

(d) Engineering constraints

- **Data bandwidth** – The raw projection stream ($\approx 3.9 \times 10^8$ bits) demands high-throughput acquisition hardware.
- **Storage** – Persistent storage of full projections imposes a volume proportional to B_{aux} .
- **Compute** – Reconstruction algorithms must process 496 projections; GPU-accelerated Krylov solvers ([4]) or PINNs ([7]) can mitigate runtime but increase memory footprints.
- **Energy** – Landauer bound sets a lower limit; actual power consumption will be orders of magnitude higher due to hardware inefficiencies.
- **Quantum control** – Implementing a generic 50-qubit unitary inversion would require specifying $> 10^{31}$ bits of control, far beyond any realistic control hardware.

(e) Projected cost and performance bounds (projection)

Assuming a storage cost of $\backslash0.02\text{perGB and a projection dataset size of }B_{\text{aux}}=3.89\times10^8\text{bits}\approx48.6\text{MB}$, the raw data storage cost is projected at $\backslash0.001\backslash$. The minimal thermodynamic energy $E_{rec}^{CT} = 1.12 \times 10^{-12}$ J is negligible compared with typical CT scanner power draw ($\approx 10^3$ W), but scales linearly with data size and thus becomes significant for larger volumes.

(f) Risks and misuse

If the framework is misapplied to justify discarding data under the premise that “knowledge exceeds information,” critical diagnostic information could be lost in medical contexts. Careful validation of any compression or measurement-reduction strategy is essential.

8. Conclusion

We have introduced a quantitative framework that compares the algorithmic information required to reverse a scrambling operation with the intrinsic Shannon entropy of the scrambled system. Applying the framework to a realistic 3-D cone-beam CT phantom reveals that the auxiliary data needed for exact reversal can exceed the system's own information by more than two orders of magnitude, and that the associated Landauer energy, while small, scales with the excess knowledge. A quantum analysis shows that a generic 50-qubit unitary demands description length exceeding the system's entropy by twenty-nine orders of magnitude, incurring a Landauer cost of 3.6×10^{10} J. These results expose a concrete thermodynamic and logical tension with the assumption that unitarity guarantees recoverability without external cost. The analysis also highlights conditions under which the paradox can be avoided—namely, when reconstruction algorithms exploit strong priors or compress measurement data effectively. Future work should experimentally probe the boundary between feasible and infeasible reconstruction, extend the analysis to deeper quantum scrambling, and explore cost-optimal designs for imaging and quantum-error-correction systems.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.43-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by openai/gpt-oss-120b, zai-org/glm-5.3-flash, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Towards effective information content assessment: analytical derivation of information loss in the reconstruction of random fields with model uncertainty. arXiv:2305.13870v1. <https://arxiv.org/abs/2305.13870v1> [2] Reconstruction of Binary Functions and Shapes from Incomplete Frequency Information. arXiv:1104.0446v3. <https://arxiv.org/abs/1104.0446v3> [3] Physics-informed self-supervised deep learning reconstruction for accelerated first-pass perfusion cardiac MRI. arXiv:2301.02033v1. <https://arxiv.org/abs/2301.02033v1> [4] Software Implementation of the Krylov Methods Based Reconstruction for the 3D Cone Beam CT Operator. arXiv:2110.13526v1. <https://arxiv.org/abs/2110.13526v1> [5] Tabletop Roleplaying Games as Procedural Content Generators. arXiv:2007.06108v2. <https://arxiv.org/abs/2007.06108v2> [6] Coding With Action-dependent Side Information and Additional Reconstruction Requirements. arXiv:1202.1484v1. <https://arxiv.org/abs/1202.1484v1> [7]

Cosmo-PINN: A Physics-Informed Neural Network for Cosmological Reconstruction. arXiv:2605.30139v3. <https://arxiv.org/abs/2605.30139v3> [8] Live Reconstruction of Large-Scale Dynamic Outdoor Worlds. arXiv:1903.06708v2. <https://arxiv.org/abs/1903.06708v2>

Appendix A. Divergence report

No divergent claims were identified among the drafts; all substantive statements converged on the same quantitative derivations and literature citations.

Appendix B. Claim attribution

Claim ID	Statement	Source Draft(s)	Agreement
C1	Reconstruction knowledge can exceed system information by a factor of ~ 114 for a 3-D CT phantom.	A, B	Convergent
C2	Landauer bound for erasing the auxiliary data at 300 K is 1.12×10^{-12} J for the CT case.	A, B	Convergent
C3	Binary signal of length 400 can be reconstructed from 33 Fourier coefficients (8.25%).	A, B	Convergent
C4	Krylov CGLS converges faster than PSIRT on the $256 \times 256 \times 52$ phantom.	A, B	Convergent
C5	Procedural content generation can produce complex structures from compact rule sets.	A, B	Convergent
C6	Action-dependent side information introduces a cost-constrained trade-off in reconstruction.	A, B	Convergent
C7	Cosmo-PINN reconstructs $w_{DE}(z)$ from late-time data by embedding physics in the loss.	A, B	Convergent
C8	Live dynamic reconstruction uses per-object volumes to avoid ghost artifacts.	A, B	Convergent
C9	Generic 50-qubit unitary description length $\approx 1.26 \times 10^{31}$ bits, ratio $\approx 2.5 \times 10^{29}$, Landauer cost $\approx 3.6 \times 10^{10}$ J.	B	Convergent (unique to B, retained)