

What Compressed Measurement Statistics Preserve: Fixed-Space Inverse Problems and a Walnut CT Counterexample

Daniel Schmitter

Abstract

Quadratic measurement statistics permit later changes of prior without retaining observations, provided the represented measurement space stays fixed. We give the objective identity, information equivalence of fitted-model messages, and a null-space obstruction to recovering newly enriched features. A real walnut CT slice tests compression using cardinal and cosine spaces of equal dimension and quantized-raw controls. Every cardinal point is dominated by its matched cosine point; eight-bit raw measurements with nonnegative total variation are smaller and more accurate than the two larger cardinal messages. A matched-prior diagnostic preserves that ranking within its finite solver budget. The positive result is accurate fixed-space reuse, not a unique spline information advantage or defect-inspection capability. The contribution is a precise retention boundary with complete negative measurement-prediction evidence.

1 Introduction

A distributed inspection system might retain compact evidence and answer a new question later. Changing the prior inside an existing representation and introducing genuinely new measurement features are different operations. The former can be exact for least squares; the latter generally requires additional information. We analyze this distinction and test it on measured tomography.

2 Related work and measurement model

The walnut data and finite-domain fan-beam matrix come from Hämäläinen and colleagues [1]. Coefficient-domain integrals are classical spline calculus [2]. The nonlinear-prior diagnostic uses established ADMM [3], not a new optimizer. Let A map pixel averages to projections and B map coefficients to pixel averages, so $H=AB$. Exact spline pixel averages do not make the supplied pixel operator an exact continuous ray integral.

3 Architecture and information flow


The communication boundary is between acquisition and later reconstruction. Donors use a known measurement model to form statistics; the recipient chooses a prior and solves in the same representation. Shared operators and cached normals are not transmitted in every message, but remain necessary system state. The raw-data control crosses a different boundary: it sends quantized observations and reconstructs with an ordinary image prior.

The counterexample concerns enrichment of information, not merely a denser coefficient array. A fine spline representation can contain the old image exactly while adding directions whose measured correlations were never retained. Those correlations cannot be recovered by refinement. A known invertible fitted-model message can, however, encode the same old right-hand side; it is therefore a stronger comparator than averaging donor predictions.

4 Fixed-space retention

For fixed donor operators, retain their normal matrices and right-hand sides. Expanding the squared residuals gives, for any later coefficient prior or extended-valued constraint R ,

What crosses the measurement boundary?



Acquisition above; later reconstruction or query below. Shared state is part of the system.

Figure 1: What crosses the measurement boundary?. Boxes distinguish supplied information, fitted components, and the quantity evaluated. Arrows show computation or data dependence, not a newly trained deep network.

$$\frac{1}{2} \sum_i \|H_i c - y_i\|^2 + R(c) = \frac{1}{2} c^T J c - b^T c + R(c) + \frac{1}{2} \sum_i \|y_i\|^2, \quad (1)$$

$$J = \sum_i H_i^T H_i, \quad b = \sum_i H_i^T y_i. \quad (2)$$

Proposition 1. J and b determine the minimizers for any such R over the fixed space. Absolute objective values also need the final constant. The proof is the expansion. Changing observation weights, forward calibration, or a nonlinear noise model need not preserve sufficiency. Raw-data removal also does not imply privacy.

For nonnegative pixel total variation, the recipient accepts only J , b , B , and the prior weight. ADMM splits $z = Bc$ and $v = DBc$, yielding

$$[J + \rho B^T B + \rho B^T D^T D B]c = b + \rho B^T (z - u) + \rho B^T D^T (v - t). \quad (3)$$

Nonnegative projection updates z and isotropic shrinkage updates v . D is the discrete pixel gradient; this is not exact continuous TV. Tensor structure compiles auxiliary normals, but the empirical measurement Gram is not assumed circulant.

5 Equivalent messages and missing features

If $J_i + \lambda Q$ is known and invertible, a fitted model $c_i = (J_i + \lambda Q)^{-1} b_i$ determines b_i by multiplication. Thus model coefficients can preserve the same information as a right-hand-side message. Naively averaging donor predictions is not an adequate control for a unique evidence-retention claim.

Proposition 2. Old right-hand sides cannot generally reconstruct enriched features. If h is outside the column space of H_0 , let $r = (I - P)h$, where P projects onto that space. Then

$$H_0^T (y + r) = H_0^T y, \quad (4)$$

$$h^T (y + r) - h^T y = \|r\|^2 > 0. \quad (5)$$

Identical old statistics require different new statistics, so no deterministic decoder is correct for both. This is a linear-algebra counterexample, not an assertion that the perturbation is a physically realizable specimen. If $H_1 = H_0^T$ instead, exact transport gives $b_1 = T^T b_0$ and $J_1 = T^T J_0 T$. Function-preserving spline refinement does not place all new fine-space directions in the old measurement span.

6 Experimental methods

The supplied operator has 9,840 rows and 6,724 columns: 120 views of an 82-by-82 slice. Zero-based views not congruent to three modulo four give ninety acquisition views, thirty per donor. Views congruent to three modulo eight give fifteen development-validation views. The fifteen views congruent to seven modulo eight remain unscored.

Cardinal cubic and cosine spaces have equal dimensions 324, 676, and 1,156. Separable pixel averages form their field maps. The initial comparison uses continuous gradient-energy quadratic priors; raw binary32, sixteen-bit, and eight-bit messages receive quadratic and nonnegative TV controls. Every fixed regularization trial is retained, with validation selection. The stronger raw TV prior motivates a separate matched-prior diagnostic.

That diagnostic gives all six compact spaces the same five TV weights as the raw control. Selected fits receive fixed 6,000-iteration checks at tolerance 10^{-7} ; previously selected raw fits extend to 7,200 iterations without new selection. Errors concern held measured views, not ground-truth images or defect labels. Numeric payload includes donor metadata but excludes file-container overhead and shared operators.

7 Results

Table 1: Initial admission frontier. Compact rows show cardinal / cosine errors at matched dimension.

Dimension / control	Message bytes	Validation relative RMSE
324	4,116	0.108249 / 0.092104
676	8,340	0.078360 / 0.069532
1,156	14,100	0.063146 / 0.055151
Raw 8-bit + TV	7,608	0.042262
Raw 16-bit + TV	14,988	0.042184
Raw binary32 + TV	29,748	0.042184

Every cardinal point is dominated by matched cosine. Eight-bit raw is smaller and more accurate than the two larger cardinal messages. Known-matrix decoding of coefficient messages reproduces the largest cardinal error within approximately 3×10^{-8} absolute difference, consistent with information equivalence.

Tighter TV checks yield cardinal errors 0.117391, 0.084271, and 0.062097, versus cosine 0.101446, 0.074864, and 0.055911. All six miss strict primal feasibility and retain small negative pixels: these are finite-budget results, not constrained optima. Extended raw errors are 0.04224944 for binary32 and 0.04230795 for eight-bit. The tested prior does not reverse the ranking.

Binary32-retained versus full-statistic reconstructions agree within 2.13×10^{-6} relative image norm for cardinal and 3.30×10^{-8} for cosine. The executable admission receipt counts 11,683,492 shared operator bytes including sparse indices and pointers; older narrative text reported 11,671,028. Additional normals occupy 3.36-42.76 MB, much more than the messages. Admission and diagnostic take 30.24 and 40.00 seconds, with 442.69 and 372.58 MiB peak RSS.

8 Discussion

The positive result is fixed-space objective reuse, supported by both bases and information-equivalent model messages. The attempted cardinal compression route fails. One specimen and development-view selection do not establish general imaging performance. No reserved view is opened for this reconstruction.

An inspection claim would require a costly physical decision, defect or geometry ground truth, multiple specimens, and compressed-raw reconstruction with classical targeted-acquisition controls. New measurements or retained extra statistics must pay for directions absent from the original space. The current algebra cannot manufacture discarded information.

9 Application boundary and research implication

The application opportunity is deferred reconstruction within an explicitly retained measurement family. The walnut experiment establishes that capability but rejects the tested spline compression frontier. A

later inspection system would need a new decision label and acquisition benefit, not a visually appealing reconstruction promoted into a defect-detection result.

10 Conclusion

Quadratic statistics preserve a defined family of future inverse problems, not arbitrary future features. The real-data study validates fixed-space reuse while rejecting the tested cardinal compression mechanism. Both findings delimit what an evidence-memory interface can honestly promise.

References

- [1] K. Hämäläinen et al. Tomographic X-ray Data of a Walnut. 2015. <https://arxiv.org/abs/1502.04064>
- [2] A. Badoual, D. Schmitter, and M. Unser. An Inner-Product Calculus for Periodic Functions and Curves. IEEE Signal Processing Letters 23(6), 878-882, 2016. <https://doi.org/10.1109/LSP.2016.2555139>
- [3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. Foundations and Trends in Machine Learning 3(1), 1-122, 2011. https://web.stanford.edu/~boyd/papers/admm_distr_stats.html