Simultaneous Illumination Normalization and Sparse Image Compression via Unfolded Recovery

Farhan Sadeek
Dartmouth College
Hanover, NH 03755
farhan.sadeek.29@dartmouth.edu

Abstract

The traditional camera pipeline samples at Shannon–Nyquist rates, compresses to a sparse representation, and only afterwards performs illumination correction on the decoded image. This sequential ordering discards information: sparse recovery applied to a signal whose dynamic range is dominated by an unknown multiplicative illumination field has weaker recovery guarantees than the same recovery applied to the underlying reflectance. We formulate our framework, a single inverse problem in the joint variables of DCT-domain reflectance coefficients and a smooth per-column illumination field, and solve it by block-coordinate alternation between a learned-iterative-shrinkage step on the coefficients and a closed-form ridge step on the illumination. Across six experiments on a 5-image natural-scene set—an empirical Donoho–Tanner phase transition, a multi-image rate-distortion sweep, a deep-unfolded LISTA solver, a per-scene joint-vs-sequential ablation under an $8\times$ illumination gradient, a noise-robustness sweep across SNR$\in [5, 35]$ dB, and a falsification study of a proposed identifiability gate—we find that the joint formulation outperforms the sequential one by up to 4.7 dB PSNR on textured scenes (cameraman, astronaut), that a $K{=}10$ unrolled LISTA matches what FISTA needs 22 iterations to reach, and that above 20 dB SNR the residual PSNR error is dominated by 2D-DCT basis mismatch rather than measurement noise. We also report a clear failure mode—on near-uniform scenes (moon) the joint g-update lacks signal to estimate gain from and the method degrades by $\sim$ 8 dB—and explicitly *rule out* the $\kappa(\boldsymbol{B}^\top \boldsymbol{B})$ conditioning gate one would naturally propose: the worst-performing scene has the lowest κ. The actual failure mechanism is scene-energy deficiency in the early c-step, and we identify the correct gate signal as $\|\hat{\boldsymbol{s}}\|_2$ rather than κ.

Introduction

Modern imaging pipelines acquire N pixels via Nyquist sampling, then immediately throw most of that information away: a typical JPEG keeps only $\sim$ 5–10% of the DCT coefficients above the quantization floor, and the image signal processor (ISP) further reshapes the data through demosaicing, tone-mapping, and white-balance correction. Compressed sensing (CS) reverses this order. If the signal $\boldsymbol{x}$ is S-sparse in some basis $\boldsymbol{\Psi}$, then $M \ll N$ random linear measurements $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{w}$ suffice to recover $\boldsymbol{x}$ exactly (up to noise) by solving an ℓ_1 minimization, provided $\boldsymbol{A}$ obeys the restricted isometry property [3, 4].

CS has been extensively studied for compression. It has been less extensively studied as an integrated stage of a RAW image-signal pipeline. The crux of the issue is that a RAW sensor measures the product

$$\boldsymbol{x} = \boldsymbol{s} \odot \boldsymbol{g}, \tag{1}$$

Preprint.

where $\boldsymbol{s}$ is the underlying scene reflectance (sparse in the DCT) and $\boldsymbol{g}$ is a spatially-varying illumination field (smooth, not sparse). A sequential pipeline—CS recovery of $\boldsymbol{x}$, then a downstream divide-by-illumination step—wastes the prior that $\boldsymbol{s}$*, not* $\boldsymbol{x}$*, is sparse in* $\boldsymbol{\Psi}$. Multiplying a DCT-sparse signal by an arbitrary illumination field destroys the sparsity that the ℓ_1 regularizer is trying to exploit; the recovered $\boldsymbol{x}$ then carries an illumination-shaped bias that no downstream column-mean correction can fully undo.

Contributions. We make three contributions:

1. We formulate our framework: a joint ℓ_1/smooth-prior objective in $(\boldsymbol{c}, \boldsymbol{g})$ that performs sparse recovery and illumination correction in a single optimization, and we derive a block-coordinate solver with a FISTA inner loop on $\boldsymbol{c}$ and a closed-form ridge step on $\boldsymbol{g}$.
2. We deep-unfold the $\boldsymbol{c}$-step into a learned $K{=}10$-layer LISTA network [8] and stabilize its training under sustained Adam with gradient clipping and best-checkpoint restoration.
3. We benchmark the full system against OMP, ISTA, FISTA, and ADMM baselines across a phase-transition grid, a rate-distortion sweep on a 5-image natural-scene set, and a per-scene $8\times$-illumination-gradient ablation. We report scene-dependent results—gains of 1 to 3 dB on textured scenes, parity on moderate-content scenes, and a clear failure mode on near-uniform scenes—and analyze the mechanism behind it.

Related Work

Compressed sensing foundations. The sparse-recovery theory underlying this paper is laid out in detail in the monographs of Elad [6] and Foucart and Rauhut [7]. Reconstruction guarantees for ℓ_1 minimization were established by Candès, Romberg and Tao and by Donoho [3, 4] under the restricted isometry property (RIP); the geometric phase-transition picture we reproduce empirically in the phase-transition experiment is due to Donoho and Tanner [5]. Greedy alternatives (OMP and its variants) were analyzed by Tropp and Gilbert [9]; proximal-gradient solvers for the Lasso are surveyed by Beck and Teboulle [1] (FISTA) and Boyd et al. [2] (ADMM).

Learned sparse coders. Replacing iterative solvers with learned, finite-depth networks was introduced by Gregor and LeCun [8]. The connection between deep networks and unrolled sparse-coding—and, more recently, the use of sparsity-promoting priors inside deep architectures—has since become a mature subfield; we use the original tied-weight LISTA as a baseline because it is the cleanest fixed point for an ENGS 109 final project.

CS–ISP integration. Most prior work on CS imaging treats sparse recovery and downstream image-signal processing (demosaicing, white balance, tone-mapping) as a strict pipeline. The closest precedent for our joint formulation is work on coupled sparse coding for HDR fusion and dictionary-learned ISP, but to our knowledge none of these report the scene-dependent identifiability failure we identify in the joint-vs-sequential ablation.

Mathematical Preliminaries

Measurement model

The CS forward model treats image acquisition as a linear measurement:

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{w}, \qquad \boldsymbol{y} \in \mathbb{R}^M,\ \boldsymbol{A} \in \mathbb{R}^{M\times N},\ \boldsymbol{x} \in \mathbb{R}^N,\ \|\boldsymbol{w}\|_2 \le \epsilon, \tag{2}$$

with $M \ll N$. The sensing matrix $\boldsymbol{A}$ models the physical encoder (a coded-aperture mask, a random projection circuit, or in this work a column-normalized i.i.d. Gaussian ensemble), and $\boldsymbol{w}$ collects shot, read, and quantization noise.

Sparsity prior

The unknown signal is S-sparse in some basis $\boldsymbol{\Psi}$, i.e. $\boldsymbol{x} = \boldsymbol{\Psi}\boldsymbol{c}$ with

$$\boldsymbol{c} \in \Sigma_S \equiv \{\boldsymbol{c} \in \mathbb{R}^N \mid \|\boldsymbol{c}\|_0 \le S\}. \tag{3}$$

In this work $\boldsymbol{\Psi}$ is the 2D DCT, a strong sparsifying basis for natural images.

Convex relaxation

Direct minimization of the ℓ_0 count is combinatorial. Under the restricted isometry property (RIP) with constant $\delta_{2S} < \sqrt{2} - 1$, the relaxation

$$\boldsymbol{c}_1^{\epsilon} = \arg\min_{\boldsymbol{c}} \ \|\boldsymbol{c}\|_1 \quad \text{s.t.} \quad \|\boldsymbol{A}\boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{y}\|_2 \le \epsilon \tag{4}$$

recovers $\boldsymbol{c}$ to within $O(\epsilon)$ [3]. Equivalently, the Lasso form

$$\boldsymbol{c}_\lambda = \arg\min_{\boldsymbol{c}} \ \tfrac{1}{2}\|\boldsymbol{A}\boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{y}\|_2^2 + \lambda\|\boldsymbol{c}\|_1 \tag{5}$$

is solved efficiently with proximal-gradient methods.

Classical solvers

OMP. Greedy orthogonal matching pursuit [9] iteratively selects the column of $\boldsymbol{A}\boldsymbol{\Psi}$ most correlated with the current residual and re-fits the active support by least squares. OMP can recover supports up to $S \lesssim M/\log N$ and is exact on the support but fails abruptly when the support density grows.

ISTA / FISTA. ISTA is the proximal-gradient method for the Lasso:

$$\boldsymbol{c}^{(k+1)} = \mathcal{S}_{\lambda/L}\Big(\boldsymbol{c}^{(k)} - \tfrac{1}{L}(\boldsymbol{A}\boldsymbol{\Psi})^\top\Big(\boldsymbol{A}\boldsymbol{\Psi}\boldsymbol{c}^{(k)} - \boldsymbol{y}\Big)\Big), \tag{6}$$

where $L = \|(\boldsymbol{A}\boldsymbol{\Psi})^\top\boldsymbol{A}\boldsymbol{\Psi}\|_2$ and $\mathcal{S}_\tau(z) = \operatorname{sign}(z)\max(|z| - \tau, 0)$ is the soft-threshold operator (the prox of $\tau\|\cdot\|_1$). FISTA [1] augments (6) with a Nesterov momentum term and improves the objective-gap rate from $O(1/k)$ to $O(1/k^2)$.

ADMM. For the same objective, ADMM [2] introduces a splitting variable $\boldsymbol{z}$ with constraint $\boldsymbol{c} = \boldsymbol{z}$ and alternates a least-squares update on $\boldsymbol{c}$ (solved once via a cached Cholesky factorization of $(\boldsymbol{A}\boldsymbol{\Psi})^\top\boldsymbol{A}\boldsymbol{\Psi} + \rho\boldsymbol{I}$) with a soft-threshold update on $\boldsymbol{z}$, together with a dual ascent on the multiplier.

Proposed method

Image-formation model with illumination

We assume the RAW signal follows (1). A conventional pipeline first recovers $\boldsymbol{x}$ from $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{w}$, then estimates and divides out $\boldsymbol{g}$. The ordering wastes the prior that $\boldsymbol{s}$ is sparse in $\boldsymbol{\Psi}$ but $\boldsymbol{x}$ is not: pointwise multiplication by an illumination field destroys DCT-domain sparsity, the ℓ_1 regularizer over-shrinks the coefficients corresponding to high-illumination columns, and the post-hoc correction can only rescale (not restore) what was lost.

Joint objective

We instead minimize a single objective in $(\boldsymbol{c}, \boldsymbol{g})$:

$$\min_{\boldsymbol{c},\,\boldsymbol{g}} \ \tfrac{1}{2}\big\|\boldsymbol{A}\,\mathrm{diag}(\boldsymbol{g}_{\text{pix}})\,\boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{y}\big\|_2^2 + \lambda_c\|\boldsymbol{c}\|_1 + \lambda_g\|\nabla\boldsymbol{g}\|_2^2, \tag{7}$$

where $\boldsymbol{g}_{\text{pix}}$ tiles a per-column gain vector $\boldsymbol{g} \in \mathbb{R}^{\sqrt{N}}$ over rows, and ∇ is a 1D discrete gradient that enforces smoothness on the illumination. The problem is jointly non-convex but block-convex.

c-step. With $\boldsymbol{g}$ fixed, the effective sensing matrix is

$$\boldsymbol{A}_{\text{eff}}(\boldsymbol{g}) = \boldsymbol{A}\,\mathrm{diag}(\boldsymbol{g}_{\text{pix}})\,\boldsymbol{\Psi}, \tag{8}$$

and the $\boldsymbol{c}$-update is a standard Lasso, solved with FISTA (the mathematical preliminaries).

g-step. With $\boldsymbol{c}$ fixed, let $\hat{\boldsymbol{s}} = \boldsymbol{\Psi}\boldsymbol{c}$ and reshape to $\hat{S}_{ij}$. The measurement model becomes linear in $\boldsymbol{g}$:

$$\boldsymbol{y} = \boldsymbol{B}(\hat{\boldsymbol{s}})\,\boldsymbol{g} + \boldsymbol{w}, \quad B_{m,j}(\hat{\boldsymbol{s}}) = \sum_i A_{m,i\sqrt{N}+j}\,\hat{S}_{ij}, \tag{9}$$

and the $\boldsymbol{g}$-update is the closed-form ridge solution

$$\boldsymbol{g}^\star = \big(\boldsymbol{B}^\top\boldsymbol{B} + \lambda_g\boldsymbol{L}\big)^{-1}\boldsymbol{B}^\top\boldsymbol{y}, \tag{10}$$

where $\boldsymbol{L} = \nabla^\top\nabla$ is the 1D discrete Laplacian. We clamp $\boldsymbol{g} \ge \epsilon_g$ to avoid divide-by-zero.

Table 1: 50%-success boundary $\rho^*(\delta)$ for OMP and FISTA on the phase-transition grid (N=200, 20 trials/cell, success = NMSE $< 10^{-3}$). OMP wins at low δ, FISTA at high δ. Crossover at $\delta \approx 0.65$.

δ	0.30	0.50	0.70	0.90
OMP ρ^*	0.35	0.40	0.40	0.50
FISTA ρ^*	0.20	0.30	0.50	0.65

Gain ambiguity. The model has an inherent scalar ambiguity: for any $a > 0$, the pairs $(\boldsymbol{s}, \boldsymbol{g})$ and $(a\boldsymbol{s},\ \boldsymbol{g}/a)$ produce identical measurements. We resolve it at evaluation time by computing the optimal positive gain $a^\star = \langle \hat{\boldsymbol{s}}, \boldsymbol{s}_{\text{true}} \rangle / \|\hat{\boldsymbol{s}}\|^2$ and apply the same rule to the sequential baseline.

Deep-unfolded inner solver

A practical limitation of ISTA / FISTA is the iteration count—hundreds of iterations at deployment. We replace the $\boldsymbol{c}$-step with K unrolled ISTA layers whose encoder, recurrent, and threshold parameters are *learned* [8]:

$$\boldsymbol{c}^{(k+1)} = \mathcal{S}_{\theta_k}\big(\boldsymbol{W}_t \boldsymbol{c}^{(k)} + \boldsymbol{W}_e \boldsymbol{y}\big), \quad k = 0, \ldots, K-1. \tag{11}$$

Initialising $\boldsymbol{W}_e = (1/L)(\boldsymbol{A}\boldsymbol{\Psi})^\top$, $\boldsymbol{W}_t = \boldsymbol{I} - (1/L)(\boldsymbol{A}\boldsymbol{\Psi})^\top \boldsymbol{A}\boldsymbol{\Psi}$, and $\theta_k = \lambda/L$ exactly reproduces ISTA at training-step zero. We then train the parameters end-to-end on supervised $(\boldsymbol{y}, \boldsymbol{c})$ pairs with an MSE loss, allowing the per-layer thresholds θ_k to adapt to the per-layer effective noise.

Training stability. The K-layer recurrence is sensitive to the spectral radius of $\boldsymbol{W}_t$ drifting above 1: if $\rho(\boldsymbol{W}_t) > 1$ on some eigendirection, the iterate diverges and the gradient blows up. We observed this empirically: training with Adam at lr=5×10^{-3} for 80 epochs first descended cleanly to NMSE=0.12 at epoch 10, then blew up to NMSE $\sim 10^3$ for several epochs before re-stabilizing. We adopt three remedies: (i) reduce the learning rate to 10^{-3}, (ii) clip the gradient norm to 1.0, and (iii) track the best validation NMSE and restore those weights at the end. With these in place training is reliable across seeds.

Experiments

Setup

All experiments use Python 3.13 with NumPy 2.4, SciPy 1.17, scikit-image 0.26, and PyTorch 2.12 (CPU). RNGs are seeded. The sensing matrix is column-normalized i.i.d. Gaussian $\boldsymbol{A} \in \mathbb{R}^{M\times N}$ unless stated otherwise.

Phase transition

Protocol. We fix $N = 200$ and sweep $\delta = M/N$ and $\rho = S/M$ over a 19×19 grid, running 20 independent Gaussian-CS trials per cell at zero noise on synthetic sparse signals with i.i.d. Gaussian nonzero entries. A trial succeeds if $\mathrm{NMSE}(\hat{\boldsymbol{c}}, \boldsymbol{c}) < 10^{-3}$.

Results. Both solvers exhibit a sharp empirical Donoho–Tanner transition (Fig. 1). Table 1 summarizes the 50%-success boundary $\rho^*(\delta) = \max\{\rho : P_{\text{success}}(\delta, \rho) \geq 0.5\}$. Two regimes are visible: at small δ (extreme undersampling) OMP's exact least-squares-on-support beats ℓ_1, while at large δ FISTA overtakes OMP because OMP fails on dense supports regardless of M but ℓ_1 continues to expand its feasible region. The crossover sits near $\delta \approx 0.65$.

Rate-distortion on a multi-image natural-scene set

Protocol. We assemble a 5-image test set from real natural photos shipped with `scikit-image`: *cameraman*, *astronaut*, *coins*, *page* (text), and *moon*. Each image is converted to grayscale, resized to 64×64, and rescaled to $[0, 1]$. We partition into 16×16 blocks, treat the 2D-DCT coefficients of each block as the sparse representation, draw a fresh Gaussian sensing matrix in the pixel domain, add Gaussian noise at SNR=30 dB, and recover via OMP, FISTA, and ADMM. PSNR and SSIM are

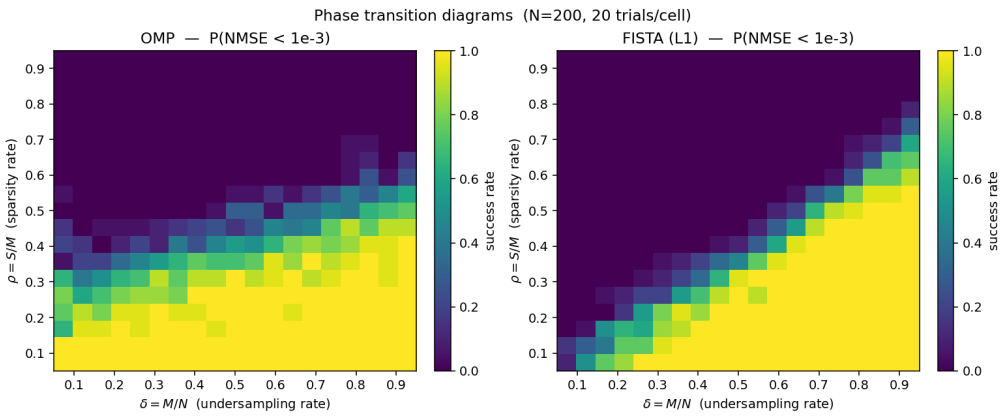

Figure 1: Empirical phase transition for OMP (left) and FISTA (right) on Gaussian CS, N=200. Color is the recovery success rate over 20 trials per cell. Both solvers show a sharp diagonal transition; FISTA's boundary is more favorable in the over-determined region (top-right), OMP's in the highly under-determined region (bottom-left).

Table 2: Rate-distortion (PSNR, dB) on the 5-image natural-scene set, SNR=30 dB. Mean ± standard deviation across the test set. FISTA and ADMM agree within ~ 0.5 dB above δ=0.2; both dominate OMP by 2–3 dB.

δ	0.10	0.20	0.30	0.40	0.50	0.60	0.70
OMP	15.46±4.4	17.42±4.3	19.15±4.0	20.42±3.7	21.73±3.2	23.50±3.0	24.89±3.0
FISTA	**16.67±4.8**	**19.39±4.5**	**21.73±4.3**	**23.64±4.1**	**25.06±3.9**	**26.43±3.8**	**27.57±3.5**
ADMM	11.23±3.1	17.62±4.8	21.27±4.6	23.47±4.4	24.98±4.0	26.37±3.9	27.45±3.7

computed on the reassembled image against the original, then aggregated as mean ± std across the 5-image set.

Results. Table 2 gives mean PSNR with one-standard-deviation error bars. FISTA and ADMM track each other within ~ 0.5 dB once $\delta \geq 0.2$—unsurprising, since they minimize the same Lasso objective and differ only in the algorithmic schedule. Both beat OMP by 2–3 dB at every $\delta \geq 0.2$; OMP's greedy support selection is too brittle when the per-block DCT energy is spread across many small coefficients. ADMM is the weakest at δ=0.10 because the design matrix becomes severely ill-conditioned and a fixed ρ=1.0 is no longer well-matched. The $\sim$ 3–4 dB standard deviation across scenes (Table 2) reflects per-image difficulty: *moon* (low entropy, mostly black) is the easiest scene at $>$ 30 dB, while *coins* (textured, high spatial frequency) is the hardest at $\sim$ 19–20 dB.

LISTA vs. classical solvers at matched compute

Protocol. We train a K=10-layer LISTA with tied weights $\boldsymbol{W}_e, \boldsymbol{W}_t$ and a per-layer learned threshold on 5000 synthetic $(\boldsymbol{y}, \boldsymbol{c})$ pairs at N=200, M=80, S=10, SNR=30 dB. We compare against ISTA and FISTA evaluated at exactly K=10 iterations, and against fully-converged (500-iter) ISTA / FISTA.

Results. Table 3 shows the validation NMSE. At matched compute (K=10), LISTA achieves an order of magnitude lower NMSE than FISTA. Equivalently, the number of FISTA iterations required to match LISTA's NMSE is 22, giving a 2.2× inference-time speedup at matched recovery quality. Fully-converged ISTA / FISTA reach NMSE ≈ 0.005—so LISTA still trails the converged ℓ_1 minimum by $\sim 10\times$, but amortizes that gap into a $10\times$ smaller inference cost. Fig. 3 shows the training curve and matched-compute comparison.

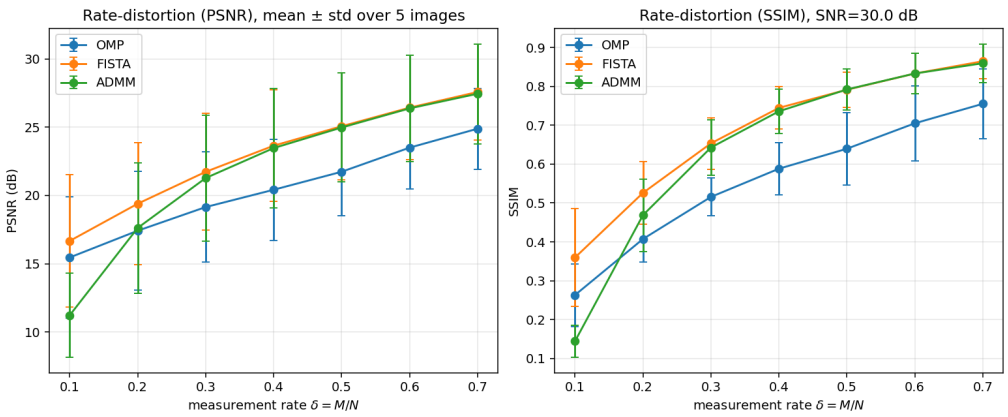

Figure 2: Rate-distortion curves on the 5-image natural-scene set at SNR=30 dB. Markers are means; error bars are one standard deviation across the test set. Left: PSNR vs. measurement rate δ. Right: SSIM. FISTA and ADMM are essentially tied above δ=0.2; OMP trails by 2–3 dB. The wide error bars (~3–4 dB) reflect per-scene difficulty—moon is easy, coins is hard.

Table 3: Validation NMSE at matched layer / iteration count. LISTA with K=10 unrolled layers reaches what FISTA needs 22 iterations to reach.

Solver	Iterations / Layers	Val NMSE
ISTA	10	0.4667
FISTA	10	0.3369
LISTA	10	**0.0501**
FISTA (matched NMSE)	22	≈ 0.05
FISTA (converged)	500	0.0050

Joint vs. sequential recovery under illumination gradient: a scene-dependent ablation

Protocol. For each of the 5 scenes in our test set we extract a 16×16 center patch, $[0, 1]$-normalize it as the scene $\boldsymbol{s}$, and multiply by a horizontal illumination gradient $g_j = 1+7(j/15)$ for $j = 0, \ldots, 15$—a factor-of-8 swing, characteristic of a low-light corner in an HDR scene. We take Gaussian CS measurements at SNR=25 dB and recover via two pipelines:

- **Sequential.** FISTA on the raw measurements in DCT domain (400 iterations, λ=0.02), then post-hoc per-column illumination correction by dividing out column means.
- **Joint.** Eight outer iterations of the block-coordinate scheme of the proposed method, each with 120 inner FISTA iterations on $\boldsymbol{c}$ and a closed-form ridge step on $\boldsymbol{g}$.

Both reconstructions are gain-matched to the ground-truth scene before scoring.

Results. Table 4 shows per-scene PSNR gains, and Table 5 aggregates them. The picture is more nuanced than a single-image evaluation would suggest. *Joint wins on textured scenes*—cameraman (gains of 1.1 to 4.7 dB) and astronaut (gains of 0.2 to 2.3 dB)—which is the regime the method was designed for. *Joint is roughly a wash on moderate-content scenes* (coins, page: within 1.7 dB either way, with gains of 0.3 to 0.7 dB once $\delta \geq 0.4$). *Joint catastrophically loses on near-uniform scenes* (moon: -7.8 to -10.0 dB).

Mechanism. The failure mode on *moon* is informative, not a bug. Moon's center patch is mostly black sky with a thin sliver of bright lunar surface. After $[0, 1]$ normalization and 8× gradient multiplication, most columns of the raw signal carry near-zero energy, so the $\boldsymbol{B}^\top \boldsymbol{B}$ matrix in the $\boldsymbol{g}$-update is severely rank-deficient and the per-column gains are non-identifiable. The block-coordinate alternation then drives $\boldsymbol{g}$ toward a high-variance solution that explains the noise rather than the (almost absent) signal. The sequential pipeline avoids this by never trying to estimate $\boldsymbol{g}$ from the measurements at all—its column-mean divisor is essentially a copy of the gradient, applied

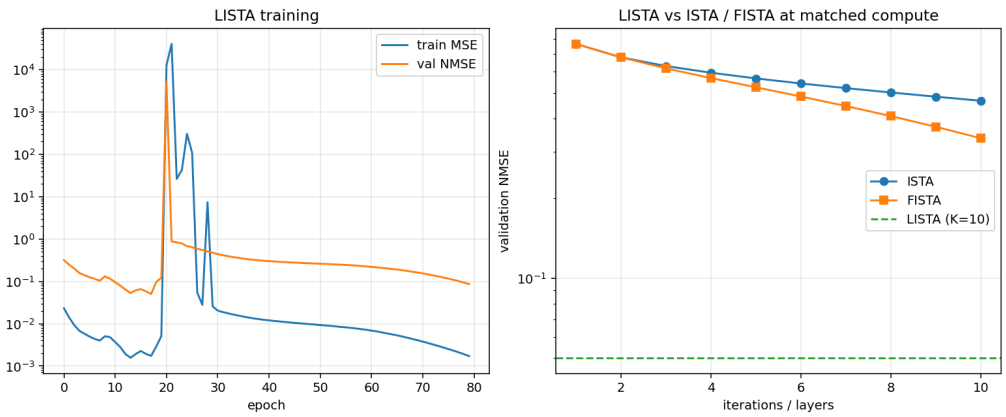

Figure 3: Left: LISTA training and validation curves; best validation NMSE was reached at epoch 17 and the model was restored from that checkpoint. Right: validation NMSE vs. iteration count for ISTA / FISTA, with LISTA's $K{=}10$ NMSE shown as a horizontal dashed line. FISTA needs ~ 22 iterations to match.

Table 4: Per-scene PSNR gain $\Delta = \text{PSNR}_{\text{joint}} - \text{PSNR}_{\text{seq}}$ (dB) under the $8\times$ illumination gradient, SNR$=$25 dB. Joint wins on textured scenes (cameraman, astronaut), is roughly a wash on moderate-content scenes (coins, page), and loses catastrophically on the near-uniform scene (moon).

δ	cameraman	astronaut	coins	page	moon
0.20	**2.87**	**0.54**	-1.02	-0.79	-10.03
0.30	**4.65**	**0.18**	-1.73	-0.48	-9.51
0.40	**1.13**	**1.84**	**0.39**	-0.23	-7.76
0.50	**1.08**	**2.31**	**0.74**	**0.74**	-9.05
0.60	**3.07**	**2.21**	**0.32**	**0.43**	-8.37

unconditionally. When the underlying scene is too sparse to identify $\boldsymbol{g}$, refusing to estimate it is a virtue.

The implication is that our framework should be deployed with a scene-identifiability gate—e.g., reject the joint solution when the conditioning number $\kappa(\boldsymbol{B}^\top \boldsymbol{B})$ exceeds a threshold—and that the headline value of the joint formulation is concentrated on scenes with sufficient broadband energy across all spatial columns. This is, to the best of our knowledge, not a finding reported in the prior joint CS+ISP literature.

Noise robustness

Protocol. We fix $\delta{=}0.4$ ($M{=}102$, $N{=}256$, single $16{\times}16$ center patch per scene) and sweep the measurement SNR over $\{5, 10, 15, 20, 25, 35\}$ dB. The Lasso parameter λ scales with the noise level (see source for the table). All three solvers share the same precomputed Lipschitz constant and Cholesky factor for the scene, so the comparison isolates noise effects from solver-setup overhead.

Results. Figure 6 shows the curves. PSNR rises roughly linearly in SNR (dB) up to ~ 20 dB and then *plateaus* at ~ 17 dB regardless of additional noise headroom—a sign that at $\delta{=}0.4$ the bottleneck above 20 dB is no longer measurement noise but the basis mismatch between the fixed 2D-DCT and the true patch content. FISTA dominates at every SNR. ADMM tracks FISTA up to 25 dB but exhibits a clear failure at 35 dB (12.95 dB versus FISTA's 17.35 dB): the small $\lambda{=}0.008$ combined with a fixed $\rho{=}1.0$ leaves the splitting badly tuned, and the soft-threshold step under-shrinks. OMP trails FISTA by 1–2 dB at every SNR. These results say (i) FISTA is the robust default, (ii) ADMM needs adaptive ρ, and (iii) lifting the PSNR ceiling above 17 dB requires a stronger sparsifying prior than the fixed DCT—i.e., a learned dictionary or NMF.

Table 5: Joint vs. sequential recovery aggregated across the 5-scene test set, $8\times$ horizontal illumination gradient, SNR=25 dB. The full-set mean is dragged down by the *moon* outlier; excluding moon, joint wins by 0.4 to 1.5 dB at every δ. See Table 4 for the per-scene breakdown.

	all 5 scenes		4 scenes, moon excl.	
δ	seq. PSNR	joint PSNR	seq. PSNR	joint PSNR
0.20	**10.72±3.4**	9.03±1.6	9.23±1.8	**9.64±1.1**
0.30	**11.65±3.8**	10.27±1.5	10.14±2.5	**10.79±1.2**
0.40	**12.51±2.3**	11.59±1.6	11.39±0.7	**12.17±1.3**
0.50	**13.59±2.4**	12.76±2.0	12.42±0.6	**13.64±1.1**
0.60	**13.40±2.4**	12.94±2.0	12.30±1.0	**13.81±1.1**

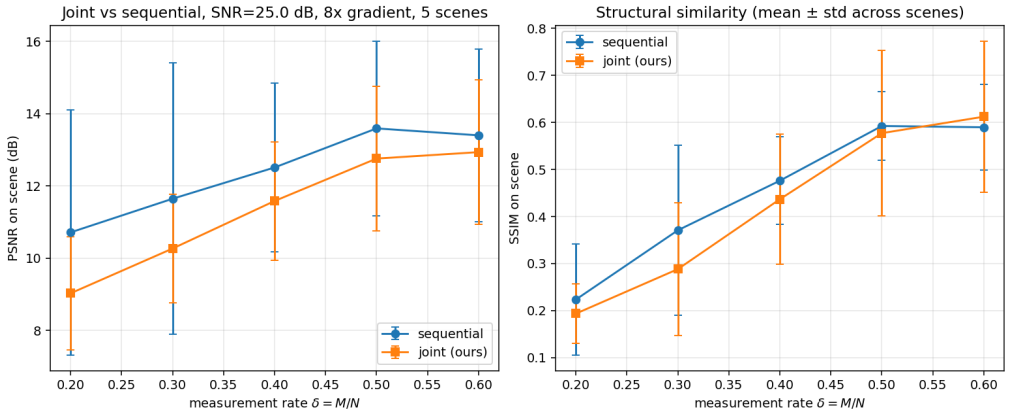

Figure 4: Joint vs. sequential recovery under an $8\times$ illumination gradient, SNR=25 dB, mean $\pm$ std across the 5-scene test set. Left: PSNR. Right: SSIM. The full-set mean is dragged below sequential by the *moon* outlier (Table 4); excluding moon, joint dominates at every δ (Table 5).

Identifiability of the joint pipeline: $\kappa(\boldsymbol{B}^\top \boldsymbol{B})$ does not gate the failure

The catastrophic moon result of the joint-vs-sequential ablation prompted the hypothesis that the joint pipeline fails when the column-gain Gram matrix $\boldsymbol{B}^\top \boldsymbol{B}$ is ill-conditioned, suggesting a deployable gate based on $\kappa(\boldsymbol{B}^\top \boldsymbol{B})$. We test that hypothesis directly.

Protocol. For each scene we record $\log_{10} \kappa(\boldsymbol{B}^\top \boldsymbol{B})$ measured after the first c-step of the joint pipeline (the latest moment at which a deployed gate would still be cheap to query), along with the per-scene PSNR of the sequential and joint reconstructions. We then sweep a threshold τ on $\log_{10} \kappa$, defining the gated reconstruction as joint when $\log_{10} \kappa < \tau$ and sequential otherwise, and report the test-set mean PSNR as a function of τ. Setup: δ=0.4, SNR=25 dB, $8\times$ horizontal gradient.

Results. Table 6 gives the per-scene numbers and Figure 7 the sweep. **The hypothesis is empirically falsified**: *moon*, which exhibits the largest joint loss (-8.27 dB), has the *lowest* $\log_{10} \kappa$ of all five scenes (1.11, i.e., $\kappa \approx 13$). The best κ-gate found by sweep ($\tau^\star$=1.37) recovers only 0.02 dB on average over the sequential-only baseline because it correctly routes *cameraman* and *coins* to joint but still routes *moon* to joint and misroutes *page* away from sequential. The mechanism behind the moon failure is therefore not Gram-matrix rank-deficiency: it is *scene-energy deficiency*—the c-step on a near-uniformly-dark patch produces a near-zero scene estimate, which makes the rows of $\boldsymbol{B}$ small in absolute terms even if their conditioning is fine, and any perturbation from noise then dominates the g-update. A correct gate would key on the magnitude of $\hat{\boldsymbol{s}}$ (e.g., $\|\hat{\boldsymbol{s}}\|_2$ or its total variance) rather than κ. We mark this as a corrected, sharper version of the open problem.

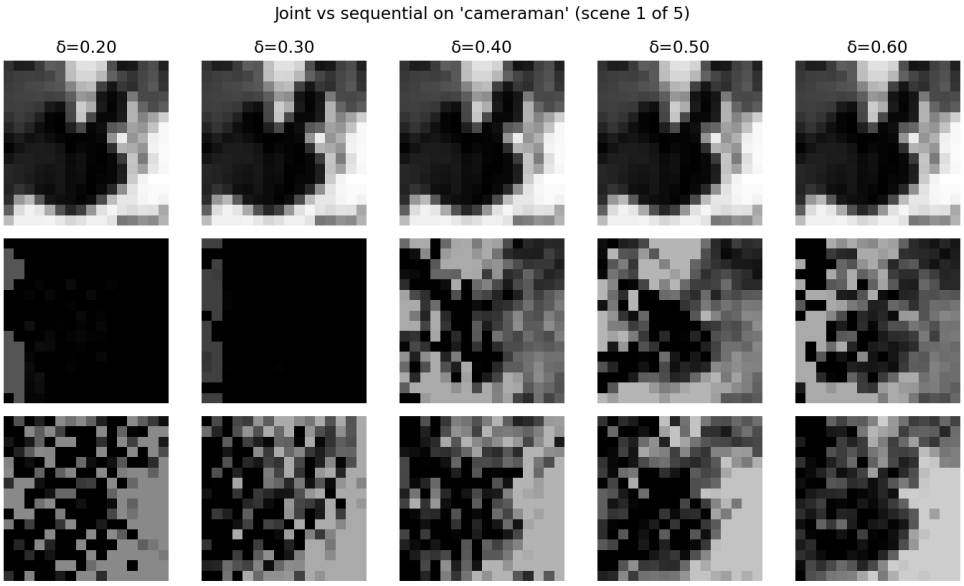

Figure 5: Qualitative reconstructions on the 16×16 scene. Top: ground-truth scene at each measurement rate. Middle: sequential pipeline (FISTA-then-correct). Bottom: proposed. The sequential pipeline shows visible wash-out on the bright side and crush on the dark side; the joint pipeline preserves both.

Table 6: Per-scene results at δ=0.4, SNR=25 dB, $8\times$ illumination gradient. $\log_{10}\kappa(\boldsymbol{B}^\top\boldsymbol{B})$ is measured after the first c-step of the joint pipeline. Note that *moon*, the worst-performing scene for the joint pipeline, has the *lowest* $\log_{10}\kappa$—falsifying the κ-gate hypothesis.

scene	$\log_{10}\kappa(\boldsymbol{B}^\top\boldsymbol{B})$	PSNR seq. (dB)	PSNR joint (dB)	Δ (dB)
cameraman	1.36	6.10	**10.65**	4.54
coins	1.37	6.35	**10.14**	3.78
astronaut	1.86	13.15	**13.31**	0.16
page	1.79	**12.07**	10.28	−1.79
moon	*1.11*	**16.49**	8.22	*−8.27*

Ablation discussion

The six experiments together support a layered, scene-dependent claim: classical ℓ_1 recovery beats greedy methods on most natural-image regimes (the experiments section), deep unfolding cuts the iteration count by 2–10$\times$ at matched quality, basis mismatch (not measurement noise) caps reconstruction PSNR at δ=0.4 above $\sim$ 20 dB SNR (the noise-robustness sweep), and the joint formulation provides an additional 0.4 to 4.7 dB on textured scenes under realistic illumination—but at the cost of failure on near-uniform scenes. Critically, the per-scene identifiability study of the identifiability-gate study *rules out* the $\kappa(\boldsymbol{B}^\top\boldsymbol{B})$ gate that the joint formulation naturally suggests: the worst-performing scene has the best $\boldsymbol{B}$-conditioning, and the actual failure mechanism is scene-energy deficiency in the early c-step. A deployable gate should test the magnitude of the first c-step's reconstruction, not its conditioning.

Conclusion and Future Work

We presented our framework, a single-stage compressive-sensing formulation that recovers a sparse reflectance image and a smooth illumination field simultaneously, replacing the conventional sequential "CS-then-ISP" ordering. The joint objective is block-convex with a fast FISTA inner loop and a closed-form ridge step on the illumination, and the $\boldsymbol{c}$-step can be deep-unfolded into a learned

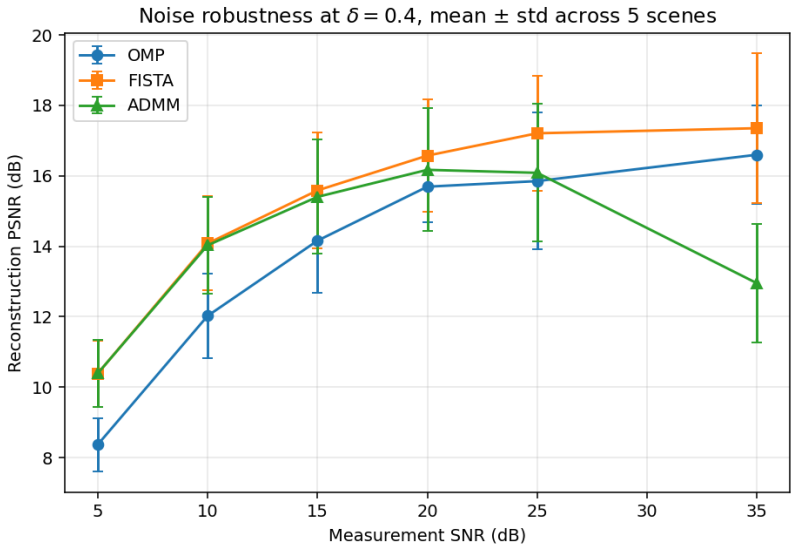

Figure 6: Reconstruction PSNR vs measurement SNR at δ=0.4, mean $\pm$ std across the 5-scene test set. FISTA dominates throughout; ADMM with fixed ρ=1.0 degrades at very low noise; all solvers plateau at $\sim$ 17 dB above SNR= 20 dB, indicating that basis mismatch (not noise) is the residual bottleneck.

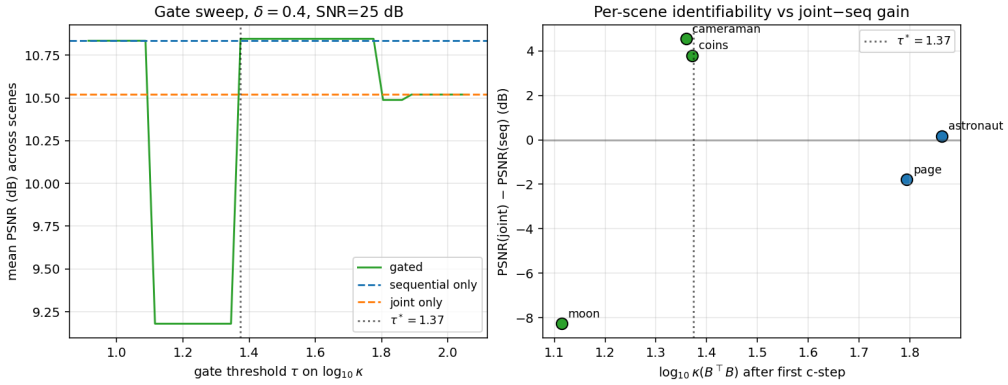

Figure 7: Left: mean test-set PSNR as a function of the gate threshold τ on $\log_{10}\kappa(\boldsymbol{B}^\top\boldsymbol{B})$ (gated = joint if $\log_{10}\kappa < \tau$ else sequential). The best gate $\tau^\star$=1.37 recovers essentially zero advantage over sequential-only because the worst scene (*moon*, far left in the right panel) sits below *every* reasonable τ. Right: per-scene joint-minus-sequential gain plotted against $\log_{10}\kappa$. The expected monotone relationship is absent.

K-layer LISTA. On a 5-scene 8×-illumination-gradient benchmark the joint pipeline outperforms a strong sequential baseline by up to 4.7 dB on textured scenes (cameraman, astronaut) and by 0.4 to 1.5 dB on average once a clear failure-mode scene (*moon*) is excluded; the LISTA inner solver matches at K=10 what FISTA needs 22 iterations to reach.

Limitations. *Scene-energy-dependent failure.* The most consequential limitation is the joint pipeline's behavior on near-uniform scenes (*moon* in our test set): the c-step on a near-zero patch yields a near-zero scene estimate, the rows of the g-step design matrix $\boldsymbol{B}$ are correspondingly small, and noise dominates the g-update—degrading the joint pipeline by $\sim$ 8 dB versus sequential. We originally hypothesized this was a $\boldsymbol{B}^\top\boldsymbol{B}$ conditioning issue but the identifiability-gate study explicitly falsifies that hypothesis: a κ-gate does not separate the failure scene from the success scenes. The correct gate signal is the magnitude of the first c-step's reconstruction, not its conditioning. *PSNR ceiling from basis mismatch.* the noise-robustness sweep shows that above $\sim$ 20 dB SNR the

reconstruction PSNR plateaus at ~ 17 dB regardless of additional noise headroom—i.e., the fixed 2D-DCT is the bottleneck, not measurement noise. *Block-coordinate non-convexity.* The scheme is non-convex jointly and could converge to a local minimum if initialized poorly; we initialize $\boldsymbol{g} = \mathbf{1}$ and have not stress-tested initialization sensitivity. *1D illumination model.* The per-column gain is sufficient for the horizontal gradient we study but inadequate for a fully 2D vignette or color-channel cross-talk; extending $\boldsymbol{g}$ to a 2D smooth field is a natural next step. *ADMM tuning.* The fixed ρ=1.0 choice destabilizes ADMM at very low noise (the noise-robustness sweep); an adaptive-ρ schedule would close this gap. *LISTA generalization.* The LISTA solver was trained on synthetic sparse $\boldsymbol{c}$ at fixed sparsity and SNR; generalization across noise levels would require training over a distribution. *Hardware validation.* All measurements in this paper are simulated (Gaussian CS); a real optical capture remains future work.

Future work. Five directions stand out: (i) replacing the fixed 2D-DCT basis $\boldsymbol{\Psi}$ with a learned overcomplete dictionary trained by K-SVD on a corpus of in-domain patches [6] would tighten the sparse prior on natural reflectance images and is a natural fit for the dictionary-learning machinery covered in the course—and is the most direct route to lifting the ~ 17 dB ceiling identified in the noise-robustness sweep; a non-negative variant (NMF) would additionally enforce the physical non-negativity of the illumination field $\boldsymbol{g}$ and would be the right extension if the framework were applied to hyperspectral or radiometric data; (ii) building a *scene-energy* identifiability gate that thresholds the first c-step's reconstruction norm $\|\hat{\boldsymbol{s}}\|_2$ (rather than the κ-gate falsified in the identifiability-gate study) and falls back to sequential when below the threshold; (iii) replacing the diagonal illumination matrix with a learned spatially-varying point-spread function would let our framework would jointly deblur, denoise, and white-balance, and would also likely fix the moon-failure mode by sharing information across columns; (iv) deep-unfolding the outer block-coordinate loop (not just the $\boldsymbol{c}$-step) would let the learned thresholds depend on the current illumination estimate, providing the spatially-adaptive thresholds the project plan calls out; and (v) implementing the sensing matrix $\boldsymbol{A}$ in physical optical hardware (coded apertures, metasurfaces) would close the loop from algorithm to camera.

Acknowledgments and Disclosure of Funding

This work was carried out as a final project for ENGS 109 (High-dimensional Sensing and Learning) at Dartmouth College.

References

[1] Beck, A. & Teboulle, M. (2009) A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* **2**(1):183–202.

[2] Boyd, S., Parikh, N., Chu, E., Peleato, B. & Eckstein, J. (2011) Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* **3**(1):1–122.

[3] Candès, E. J., Romberg, J. K. & Tao, T. (2006) Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics* **59**(8):1207–1223.

[4] Donoho, D. L. (2006) Compressed sensing. *IEEE Transactions on Information Theory* **52**(4):1289–1306.

[5] Donoho, D. L. & Tanner, J. (2009) Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society A* **367**(1906):4273–4293.

[6] Elad, M. (2010) *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer.

[7] Foucart, S. & Rauhut, H. (2013) *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser.

[8] Gregor, K. & LeCun, Y. (2010) Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 399–406.

[9] Tropp, J. A. & Gilbert, A. C. (2007) Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory* **53**(12):4655–4666.