Corruption Structure Determines Failure Mode in MRI Reconstruction

Poojak Patel* **Maneth Perera*** **Solomone Somani***
Illinois Mathematics and Science Academy

Abstract

Standard evaluation of learning-based MRI reconstruction focuses almost entirely on random undersampling, leaving open how reconstruction fails under the broader range of physically realistic measurement perturbations that arise in clinical practice. We present a controlled evaluation framework grounded in k-space physics, where the structure of any measurement corruption directly shapes reconstruction failure. Applying three physically motivated k-space corruptions to knee MRI from the fastMRI dataset (Zbontar et al., 2018), we show that corruption structure determines the qualitative character of failure independently of severity level. Random undersampling produces gradual degradation (normalized SSIM drop 0.134); structured line dropout produces catastrophic collapse with high cross-subject variance (normalized SSIM drop 0.877); and additive noise produces smooth, intermediate degradation (normalized SSIM drop 0.675). Notably, structured line dropout exhibits a high-variance regime at mild severity where visually similar corruption levels produce dramatically different failure outcomes across subjects, a non-obvious failure pattern absent from other corruption types. At severe severity, line dropout and additive noise both produce substantial degradation yet qualitatively distinct failure modes that aggregate PSNR evaluation cannot distinguish. These results argue for corruption-type-aware evaluation as a necessary component of reliable AI-based MRI reconstruction.

1 Introduction

MRI is a measurement system governed by Fourier encoding and hardware-constrained acquisition, so robustness to corruption in k-space is a physics-grounded question, not merely an image-quality question. The scanner does not directly acquire an image; it measures k-space, the two-dimensional Fourier transform of the image, via controlled magnetic field gradients (Pruessmann et al., 1999; Lustig et al., 2007). Low-frequency k-space components encode global contrast and anatomy; high-frequency components encode fine detail and edges. Any measurement corruption therefore propagates through reconstruction in ways that depend on its spatial frequency structure, not just its magnitude.

Learning-based reconstruction models have advanced rapidly on the fastMRI benchmark (Zbontar et al., 2018; Knoll et al., 2020), but fastMRI evaluates almost exclusively under random undersampling. Clinical MRI encounters a broader range of physically realistic perturbations: motion artifacts remove structured k-space bands, hardware failures cause systematic line losses, and receiver coil noise introduces additive complex Gaussian perturbations throughout k-space. Whether evaluation under random undersampling captures robustness to this broader corruption landscape is an open question.

Part of this landscape is predictable. Because low-frequency lines carry most of the image energy, removing a contiguous band near the k-space center should be more damaging than removing the

*Equal contribution.

2026 Conference on Physics and AI (PAI26).

same number of lines at random, and our aggregate severity ordering agrees; our emphasis is instead on behavior that energy arguments do not predict, namely how much the same nominal corruption level varies in effect across subjects.

We study this using zero-filled reconstruction (inverse FFT applied directly to corrupted k-space) as a deliberate baseline that introduces no learned prior, so differences in reconstruction quality across conditions can be attributed directly to corruption structure rather than model-dependent behavior. This makes the analysis a controlled characterization of measurement-space corruption rather than an evaluation of any specific architecture. The choice has a cost: because a learned reconstructor uses priors to fill in missing data, zero-filled behavior is a lower bound rather than a direct prediction of how a U-Net or variational network would fail, a point we revisit in Section 5.

Our contributions are: (1) a physics-grounded corruption evaluation framework covering 3 corruption types at 3 severity levels on the fastMRI knee dataset, with zero-filled reconstruction as a measurement-faithful baseline and total-variation (TV) regularization as a secondary check; (2) evidence that failure character, and not only severity level, depends on corruption structure, a distinction that aggregate single-corruption evaluation misses; and (3) identification of a high-variance failure regime in structured line dropout at mild severity, where similar corruption levels produce dramatically different failure outcomes across subjects, a pattern absent from the other corruption types we study.

2 Background

2.1 k-Space Physics and Reconstruction

The MRI scanner acquires complex-valued measurements in k-space whose coordinates correspond to spatial frequencies in the image (Lustig et al., 2007). This spatial frequency structure means that removing a contiguous band of k-space lines has a fundamentally different effect from removing the same number of randomly distributed lines or adding noise throughout k-space. Contiguous band removal destroys specific spatial frequency content and produces structured artifacts; random line removal aliases the image; additive noise spreads incoherently into the image domain. These qualitative differences motivate studying corruption types separately.

2.2 Learning-Based MRI Reconstruction

Deep learning models including U-Net (Ronneberger et al., 2015) and variational networks (Hammernik et al., 2018; Sriram et al., 2020) achieve strong results on fastMRI (Zbontar et al., 2018; Knoll et al., 2020). Prior work on robustness (Fang et al., 2022; Recht et al., 2020) has not provided a systematic comparison across physically distinct corruption types. Because zero-filled reconstruction introduces no learned prior, it provides a clean lower bound on reconstruction quality under each corruption condition, isolating the information content of the corrupted measurement itself.

3 Methods

3.1 Dataset

We use the fastMRI single-coil knee test set (Zbontar et al., 2018), a publicly available collection of T1-weighted and proton density fat-suppressed knee MRI acquired at 1.5T and 3T. We process 20 subjects and extract 3 central slices per subject (60 slices per condition), selecting central slices to ensure anatomically representative content across all subjects. Fully sampled k-space serves as both the reconstruction input and the source of the reference image.

3.2 Corruption Types

Random undersampling retains a protected center of 15 lines and randomly samples remaining phase-encode lines at rates of 75%, 50%, and 25% (mild, medium, severe). This models accelerated MRI acquisition and is the standard fastMRI corruption.

Line dropout zeros a contiguous block of $f \times N$ lines centered at a random offset from k-space center, with $f \in \{0.10, 0.20, 0.35\}$. Although real motion and hardware artifacts are more complex,

Table 1: PSNR and SSIM for zero-filled (ZF) and TV reconstruction ($N = 60$). $\mathrm{ND}_{\mathrm{SSIM}}$ is the SSIM drop from mild to severe. Both methods show the same failure pattern.

		Zero-Filled		TV-Regularized		$\mathrm{ND}_{\mathrm{SSIM}}$	
Corruption	Sev.	PSNR	SSIM	PSNR	SSIM	ZF	TV
Random Under.	Mild	35.97±5.4	0.931±.04	30.64	0.773		
	Med.	32.27	0.873	29.63	0.751	0.134	0.065
	Sev.	29.62	0.806	28.30	0.722		
Line Drop.	Mild	31.58	0.696	26.44	0.650		
	Med.	22.58	0.424	21.01	0.471	0.877	0.627
	Sev.	14.79	0.085	15.97	0.242		
Add. Noise	Mild	38.90	0.919	33.28	0.796		
	Med.	30.75	0.685	33.47	0.795	0.675	0.121
	Sev.	22.59	0.299	28.97	0.698		

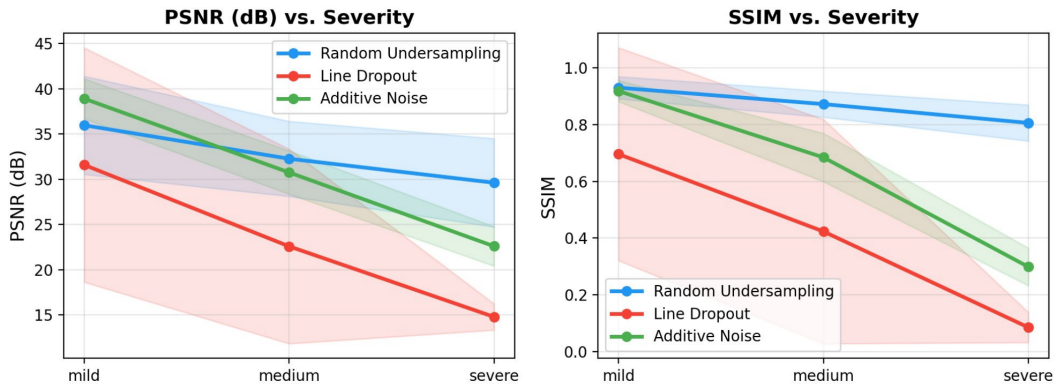

Figure 1: PSNR and SSIM vs. severity for each corruption type (shaded band = ±1 std, $N = 60$ slices). The three corruption types produce qualitatively distinct sensitivity curves.

contiguous line removal directly isolates the effect of structured missing spatial frequencies, which is the physics property of interest here.

Additive Gaussian noise adds complex-valued noise scaled to per-slice signal standard deviation at SNR levels of 20, 8, and 3 (mild, medium, severe), modeling thermal noise in MRI receiver coils.

3.3 Reconstruction and Metrics

All reconstructions use zero-filled inverse FFT on corrupted k-space, with magnitude taken as the final image. We evaluate using PSNR (dB) and SSIM (Wang et al., 2004). Normalized degradation is:

$$\mathrm{ND}(c, m) = \frac{m(c, \text{mild}) - m(c, \text{severe})}{m(c, \text{mild})} \tag{1}$$

where c is the corruption type and $m \in \{\mathrm{PSNR}, \mathrm{SSIM}\}$. To characterize the high-variance regime in line dropout, we also report cross-subject standard deviation separately and test whether line dropout variance at mild severity is significantly higher than random undersampling variance using a Levene test ($\alpha = 0.05$).

4 Results

Table 1 and Figures 1 and 2 report results. The corruption-type-dependent failure pattern is consistent across both methods, confirming it reflects measurement physics rather than reconstruction choice.

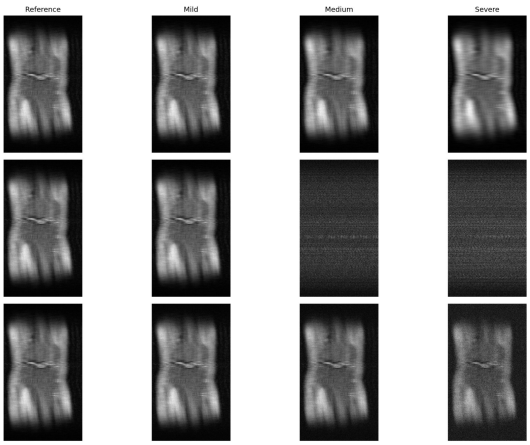

Figure 2: Representative reconstructions: rows are corruption type, columns are severity (reference at left). Row 1: random undersampling retains anatomy throughout. Row 2: line dropout collapses at medium and severe. Row 3: additive noise degrades progressively while retaining structure. All use zero-filled inverse FFT.

Random undersampling produces gradual degradation under both methods (ZF $\mathrm{ND}_{\mathrm{SSIM}} = 0.134$, TV $= 0.065$). TV over-smooths at mild severity where zero-fill is already adequate, but knee anatomy remains recognizable throughout.

Line dropout exhibits a distinctive high-variance failure regime. At mild severity, PSNR varies dramatically across subjects (31.58 ± 12.95 dB, SSIM 0.696 ± 0.375); Levene's test confirms this variance is significantly higher than random undersampling at the same level (ratio $5.7\times$, $p < 0.001$). This reflects whether the dropped band overlaps anatomically critical content, a subject-dependent outcome absent from other corruption types. At severe severity, all slices collapse under both methods (ZF SSIM 0.085, TV SSIM 0.242, $\mathrm{ND}_{\mathrm{SSIM}} = 0.877$ and 0.627 respectively).

Additive noise reveals a key reconstruction interaction. TV reduces SSIM degradation from $\mathrm{ND} = 0.675$ to 0.121 (severe SSIM 0.698 vs. 0.299), because TV denoising directly counteracts incoherent image-domain noise. This confirms that noise is recoverable in a way that structured line loss is not.

Failure character vs. failure severity. Despite both exhibiting substantial degradation at severe severity, line dropout and additive noise produce qualitatively distinct failure modes under both reconstruction methods. Line dropout yields structured banding with near-complete anatomical signal loss (ZF SSIM 0.085, TV SSIM 0.242); additive noise preserves residual structure under spatially uniform degradation (ZF SSIM 0.299, TV SSIM 0.698). The contrast is sharpest under TV, where noise is largely recoverable but structured line loss is not, indicating that failure character depends on corruption structure in a way that reconstruction strategy cannot fully compensate.

5 Discussion and Conclusion

Within the setting we study, corruption structure is a leading determinant of failure mode, separate from severity. The mild-severity line-dropout regime (variance ratio $5.7\times$, $p < 0.001$) reflects measurement physics directly: whether a dropped band is critical depends on subject-specific k-space content, a failure mode absent from standard benchmarks.

Scope and limitations. This is a preliminary characterization, not a general claim: it covers three corruption types on knee MRI with a zero-filled baseline and a single regularized method (TV), neither of which uses learned priors, so the results bound the information in each measurement rather than predicting a trained network. Testing a learned reconstructor such as the fastMRI U-Net, and enlarging the 20-subject, 3-slice cohort, are the natural next steps. The severity ordering is partly expected from k-space energy; the qualitative separation of failure modes and the high-variance line-dropout regime are not, and motivate corruption-type-aware evaluation alongside random-undersampling benchmarks.

References

Z. Fang, Y. Chen, M. Liu, L. Xia, X. Fan, Q. Hu, S. Liu, P. Liu, D. Bhatt, and S. Jain. Reliable acceleration of MRI with variable density acquisition and deep learning reconstruction. *Medical Physics*, 49(9):5882–5895, 2022.

K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6):3055–3071, 2018.

F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, et al. fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MRI technology development. *Radiology: Artificial Intelligence*, 2(1), 2020.

M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999.

M. P. Recht, J. Zbontar, D. K. Sodickson, F. Knoll, N. Yakubova, A. Sriram, T. Murrell, A. Defazio, N. Rabbani, L. Rybak, et al. Using deep learning to accelerate knee MRI at 3T: Results of an interchangeability study. *American Journal of Roentgenology*, 215(6):1421–1429, 2020.

O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015.

A. Sriram, J. Zbontar, T. Murrell, A. Defazio, C. L. Zitnick, D. Sodickson, Y. W. Lui, and F. Knoll. End-to-end variational networks for accelerated MRI reconstruction. In *MICCAI*, pages 64–73, 2020.

Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

J. Zbontar, F. Knoll, A. Sriram, T. Murrell, Z. Huang, M. J. Muckley, A. Defazio, R. Stern, P. Johnson, M. Bruno, et al. fastMRI: An open dataset and benchmarks for accelerated MRI. *arXiv preprint arXiv:1811.08839*, 2018.