Practical Advances in Artificial Intelligence and Computing
Frontier Research Publishing • https://www.frpubs.org
Vol. 4, No. 4, 2025

Equal Sampling Budgets Can Preserve Different Fine Structures in Sparse MRI Reconstruction

Ysella Vennor[1]*, Corvin Taldrake[1], Thessaly Orellis[1]
[1]University of Manitoba, Winnipeg, Canada
*Corresponding author: ysellavennor731@gmail.com

Abstract: This illustrative numerical study asks whether equal Fourier sampling budgets and exact acquired data consistency imply similar preservation of fine structure in accelerated MRI. Forty eight synthetic complex phantoms contain a fixed alternating bar region. Three sixteen row Cartesian masks are compared with zero filling and two prescribed Haar sparse reconstructions on sixty four pixel grids. Central sampling gives low global magnitude error but recovers little alternating contrast. A fixed spread mask with standard Haar shrinkage recovers mean contrast ratio 0.290 while central sampling gives 0.009, although its global error is slightly larger. Periodic sampling remains unchanged by the tested shrinkage procedure, exposing a symmetry limitation. Every final reconstruction has acquired residual below 3×10^{-16}. A constructed nonzero missing frequency perturbation leaves central acquired data unchanged. The evidence concerns a single coil Fourier proxy and does not assess patients, clinical diagnosis or trained reconstruction networks.

Keywords: accelerated MRI; Cartesian sampling; Haar sparsity; fine structure; Fourier null space; synthetic experiment

1 The acquisition question precedes the reconstruction question

Accelerated magnetic resonance imaging is often described through the number of measurements omitted from an otherwise complete acquisition. That description is operationally useful but scientifically incomplete. Two acquisitions can retain the same number of Fourier lines and remove different information about a particular structure. A reconstruction algorithm can then satisfy every acquired measurement while producing different local detail. This study examines that distinction through a small numerical experiment whose generating object, sampling masks, optimization steps and evaluation regions are fully available. The question is whether a common measurement budget and excellent acquired data consistency are sufficient to compare the preservation of a deliberately fine spatial pattern. We show that they are insufficient within the specified single coil Fourier model.

Parallel imaging established that spatial encoding can be shared between gradients and receiver sensitivity profiles. SENSE makes those sensitivities explicit in reconstruction [1], whereas SMASH synthesizes spatial harmonics through coil combinations [2] and GRAPPA estimates missing information through autocalibrating interpolation [3]. These methods explain why a measurement count alone does not define an inverse problem: the encoding operator matters. Our experiment removes coil diversity to make that dependence unusually visible. It is therefore not a comparison with clinical parallel imaging. It asks a narrower acquisition design question in a model where all methods see exactly the same complex measurements and no unreported sensitivity information can compensate for missing lines.

Compressed sensing MRI adds a second source of information through a structured representation of the image [4]. An image that is compressible in a transform domain may be recovered from fewer measurements

than its pixel count. That possibility does not mean that every equally sized subset is equally informative. Periodic omission can align with repeated image structure, while concentrated central sampling emphasizes broad intensity variation. A spread pattern can distribute ambiguity differently even if it reserves the same central region. We construct three sixteen line masks on a sixty four line grid to expose those distinctions. Their purpose is explanatory rather than to propose a universal optimal clinical trajectory.

Modern learned reconstructions combine physics with expressive priors. Variational networks [5], cascaded convolutional networks [6], and model based deep learning architectures [7] use different mechanisms to alternate information from the acquisition with learned transformations. Public benchmark efforts such as fastMRI improve the reproducibility of comparisons [8], while end to end variational networks demonstrate the value of jointly optimized components [9]. None of those models is trained or evaluated here. We instead use an explicit Haar shrinkage objective so that any difference between sampling masks can be interpreted without uncertainty about a training corpus, hidden architecture or pretrained weights.

The paper makes three linked contributions within that restricted setting. First, it separates global magnitude error from recovery of a fixed alternating bar pattern. Second, it compares three masks using identical reconstruction code, iteration count and Fourier normalization. Third, it constructs a nonzero image perturbation that produces essentially no change in the acquired central data. The last experiment is a constructive negative control, not a statistical estimate: it demonstrates the existence of an ambiguity that no acquired data residual can diagnose by itself. The resulting argument concerns information and task alignment, rather than uncertainty transport, clinical lesions or a claimed improvement over contemporary reconstruction networks.

2 A numerical object with explicitly defined local detail

The generating image has sixty four pixels along each dimension and arbitrary spatial units. Its magnitude consists of nested elliptical components, two smaller off center ellipses, and a smooth random field restricted to the main object. The outer ellipse establishes a strong boundary, the nested component creates an interior intensity difference, and the smaller ellipses prevent the image from being radially symmetric. The random field is Gaussian smoothed before addition. Every realization therefore differs slightly while retaining a common interpretable geometry. Figure 1 presents both the numerical object and an AI generated illustration of the simplified acquisition context. No panel represents a patient, a physical scanner experiment or a measured anatomical specimen.

A rectangular interior region contains alternating horizontal bars, each two pixels thick. The added bar amplitude is sampled uniformly between 0.08 and 0.16 magnitude units. The region spans a fixed range of normalized horizontal and vertical coordinates, and the exact sampled pixel masks are saved in the simulation procedure. The bars introduce a local high frequency feature whose visibility depends on phase encoding information. Because the rest of the object is unchanged in its broad geometry, an algorithm can recover the main outline and much of the global intensity distribution without accurately preserving these bars. This deliberate imbalance makes the local target distinct from an aggregate image score.

The object is complex valued. Its phase is a smooth linear combination of the two coordinates plus a small randomly weighted product term. This phase avoids granting the reconstruction a real image symmetry that would implicitly duplicate some Fourier information. The phase is not a model of a particular field inhomogeneity, flow experiment or pulse sequence. It simply ensures that the inverse problem remains a complex image problem. All reconstruction methods operate on complex arrays and are evaluated on magnitude images. The exact target magnitude is available because the image is generated numerically; that privilege is not assumed for an acquired clinical image.

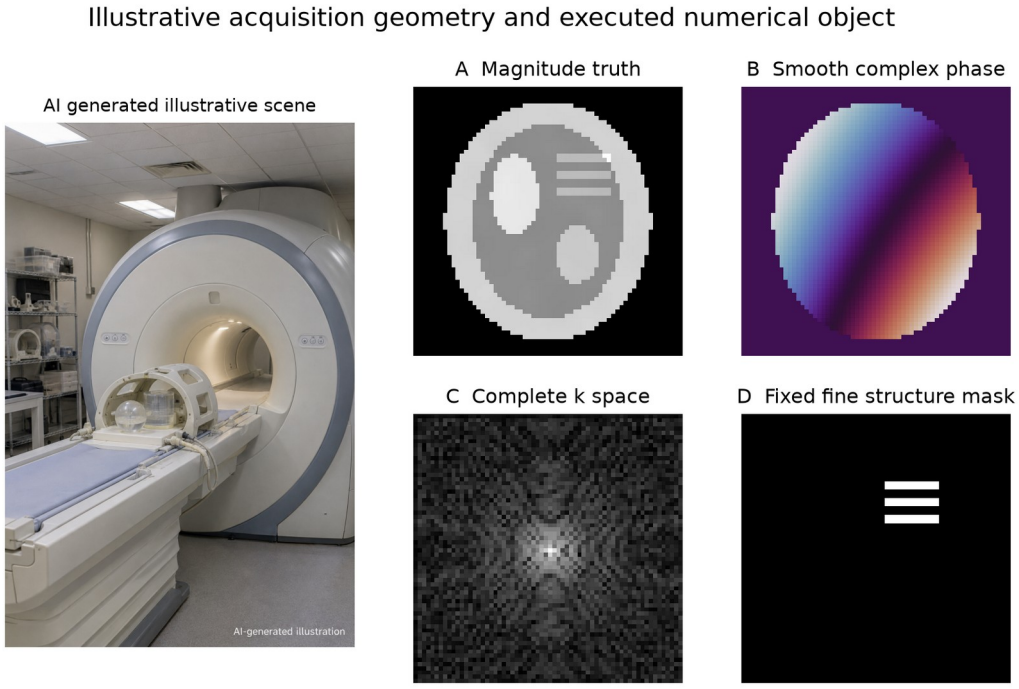


Figure 1. AI generated illustrative MRI equipment scene at left and executed digital phantom at right. The generated scene is not an actual scanner acquisition and supplies no numerical evidence. Numerical panels show magnitude truth, smooth phase, complete Fourier magnitude and fixed fine structure mask in arbitrary units.

$$x_{uv} = a_{uv} \exp(i \varphi_{uv}) \tag{1}$$

The acquisition applies a centered, orthonormal two dimensional discrete Fourier transform. A binary mask selects complete rows, corresponding to a simplified Cartesian phase encoding design. Complex Gaussian noise is added before masking, with independent real and imaginary components and a total complex standard deviation of 0.0015 in the orthonormal Fourier convention. The same complete noise array is reused across masks for a particular object. This pairing reduces irrelevant Monte Carlo variation when comparing acquisitions, but it does not make the measurements identical because each mask retains different coefficients. Noise is sampled independently between the forty eight objects.

$$y = M F x + M \varepsilon, E|\varepsilon_{uv}|^2 = \sigma^2 \tag{2}$$

Three masks retain sixteen rows each. The central mask acquires sixteen contiguous rows around zero phase encoding frequency. The periodic mask acquires every fourth row across the complete grid. The spread mask acquires eight contiguous central rows and eight rows drawn without replacement from the remaining positions. The spread draw is made once from the recorded seed and then held fixed across cases. It is consequently a comparison of three actual masks, not an average over all variable density designs. A different spread draw could change the quantitative ordering, particularly for a structure with another orientation or spatial frequency.

The measurement budget corresponds to one quarter of a fully sampled row set, but we avoid interpreting that ratio as an achieved scanner acceleration. Real acceleration depends on calibration overhead, sequence timing, preparation pulses, gradient constraints and receiver configuration. No such timing model is present. The precise claim is that all three numerical operators retain 1,024 of 4,096 Fourier coefficients. Their equal cardinality makes the comparison controlled while their different row locations create distinct null spaces.

Table 1 lists the quantities required to reproduce that statement and distinguishes parameters that are fixed from those randomized across objects.

Table 1. Executed synthetic acquisition and reconstruction specification.

Component	Specification	Interpretation
Objects	48 complex arrays of 64 × 64	No patients or acquired MRI
Sampling	16 complete rows in each mask	1,024 retained coefficients
Noise	Complex σ 0.0015	Orthonormal Fourier units
Haar solver	80 steps λ 0.0008 or 0.003	Five transform levels
Evaluation	432 ledger rows	48 objects × 3 masks × 3 methods

3 Reconstruction with an inspectable sparse prior

The first reconstruction inserts the measured coefficients into an otherwise zero Fourier array and applies the inverse transform. This zero filled estimate is a useful baseline because it adds no explicit image prior. Its artifacts directly reflect the mask and complex image. Central sampling behaves like a loss of high frequency detail, whereas periodic sampling creates coherent spatial overlap. These visual differences are not evidence that one estimator has learned to preserve a structure. They follow from the linear operator. The baseline also has essentially zero acquired coefficient residual, so it immediately shows that data consistency alone does not distinguish a detailed reconstruction from a visibly incomplete one.

$$\hat{x}_0 = F^{-1} y \tag{3}$$

The two sparse reconstructions minimize a quadratic acquired data discrepancy plus the sum of magnitudes of coefficients in an orthonormal Haar representation. The transform is implemented explicitly through separable pairwise averages and differences, recursively processing square approximation blocks of sizes sixty four, thirty two, sixteen, eight and four. We apply the same transform to complex arrays. Shrinkage therefore reduces complex coefficient magnitudes while retaining their phases, rather than thresholding real and imaginary parts independently. The complete set of transformed coefficients is penalized, including the coarsest block; no anatomy dependent mask, known bar location or generating phase is supplied to the algorithm.

$$J(x) = 0.5 \| M F x - y \|_2^2 + \lambda \| H x \|_1 \tag{4}$$

Two prescribed regularization strengths are evaluated, 0.0008 and 0.003. They are not selected using the test objects, and no validation search is performed. The lower value permits more transform detail, while the higher value imposes stronger shrinkage. Choosing two explicit values makes the regularization tradeoff inspectable without claiming that either value is optimal. The implementation uses eighty accelerated proximal gradient iterations for each reconstruction. Because the masked orthonormal Fourier operator has norm at most one, a unit gradient step is appropriate for the quadratic term. The same step size and stopping count apply to every mask and object.

$$g^k = v^k + F^{-1}(y - M F v^k) \tag{5}$$

Complex soft thresholding maps each coefficient to zero when its magnitude is below the chosen threshold and otherwise reduces its magnitude by that threshold. The inverse Haar transform then returns an image. This is a transparent instance of proximal reconstruction rather than a trained denoiser. The acceleration parameter follows the standard FISTA recurrence [10]. We do not claim a convergence certificate from a fixed eighty step run, and the script does not stop by comparing to the generating truth. A numerical optimization study could additionally report objective gaps or convergence traces, but those are separate from the acquisition comparison examined here.

$$S_{\lambda}(z)=z\,max(0,1-\lambda/|z|) \tag{6}$$

$$z^{k+1}=H^{-1}S_{\lambda}(H\,g^{k}) \tag{7}$$

After the final optimization step, the algorithm replaces every acquired Fourier coefficient by its observed value exactly. This final projection makes acquired data consistency common to all three estimators. It also means that the reported final image is a projected proximal solution rather than precisely the unprojected minimizer of Equation 4. That distinction matters because an exact projection restores the particular measured noise as well as the signal. We include the projection to make the central logical test particularly strong: even methods whose measured coefficients agree to numerical precision can disagree substantially about unmeasured fine structure.

$$\hat{x}=F^{-1}[\,M\,y+(1-M)\,F\,z^{80}\,] \tag{8}$$

Other priors would produce different reconstructions. Total variation favors sparse image gradients [11], plug and play methods insert general denoisers into iterative inversion [12], and regularization by denoising relates denoiser behavior to an explicit reconstruction framework [13]. Deep image prior exploits architectural structure without a conventional external training corpus [14]. These are methodological context rather than evaluated competitors. Their relevance is that each supplies an additional preference among images compatible with incomplete measurements. None can make the acquired operator contain a frequency it did not measure. Any claim of feature recovery must therefore combine the encoding information with assumptions about the image class.

4 Evaluation distinguishes broad fidelity from a local estimand

Global magnitude root mean squared error averages squared differences over all 4,096 pixels. It includes the background and major object boundaries as well as the small bar region. This score answers whether the reconstructed magnitude is close to the generating magnitude on average. It does not directly answer whether the alternating structure is retained. A reconstruction that smooths a small region can still obtain a favorable global score when broad intensity and geometry dominate. We therefore report a second measure using masks fixed before any reconstruction. The method never sees these masks; they are used only after reconstruction to evaluate the intended local target.

$$R=[\,N^{-2}\textstyle\sum_{u,v}(|\hat{x}_{uv}|-|x_{uv}|)^{2}\,]^{1/2} \tag{9}$$

The bar contrast is the mean magnitude in bar pixels minus the mean magnitude in the interleaved nonbar pixels inside the rectangle. The true contrast is computed from the complete generating image rather than set equal to the sampled additive amplitude. This accounts for smooth background intensity changes across alternating rows. We divide reconstructed contrast by that exact true contrast to obtain a recovery ratio. A value near one indicates preservation of this particular fixed region contrast; a value near zero indicates little alternating difference; negative values indicate a polarity reversal. This ratio is interpretable because the true contrasts are strictly positive within the generated family.

$$C(x) = mean_B|x| - mean_{R \setminus B}|x| \tag{10}$$

The acquired data residual is the norm of the difference between measured coefficients of the reconstruction and the observed coefficients, divided by the observed data norm. It measures implementation consistency with the retained noisy samples. It does not measure agreement with noise free acquired samples, nor does it test the missing coefficients. The metric is calculated after the final projection. For every method it is close to floating point precision. Reporting this small residual is useful as a numerical control but would be misleading if presented as an independent proof that the recovered bars are correct.

$$q = C(\hat{x}) / C(x), d = \| M F \hat{x} - y \|_2 / \| y \|_2 \tag{11}$$

Forty eight independent objects supply the evaluation set. The complete ledger contains one row for each object, mask and method combination, giving 432 rows. Reported mean errors and ratios are arithmetic averages across the forty eight objects in each condition. The standard deviation of contrast ratios describes variation across the specified image family; it is not a confidence interval for a clinical population. We do not use a hypothesis test to turn this explanatory simulation into a general ranking of sampling patterns. The strongest conclusions come from paired numerical changes and the constructed null space example, with the limitations of the object family kept explicit.

Figure 2 places the sampling masks beside the corresponding reconstructions of the first generated object. Its purpose is to connect the retained rows to recognizable image behavior. The selected example is fixed by generation order and is not chosen to maximize an apparent difference. All aggregate results come from the complete ledger. The ordering of images within the plate matches the ordering in the results tables, and the image display range is common across methods. This avoids giving one reconstruction a separate contrast window that could conceal amplitude loss or exaggerate fine structure visibility.

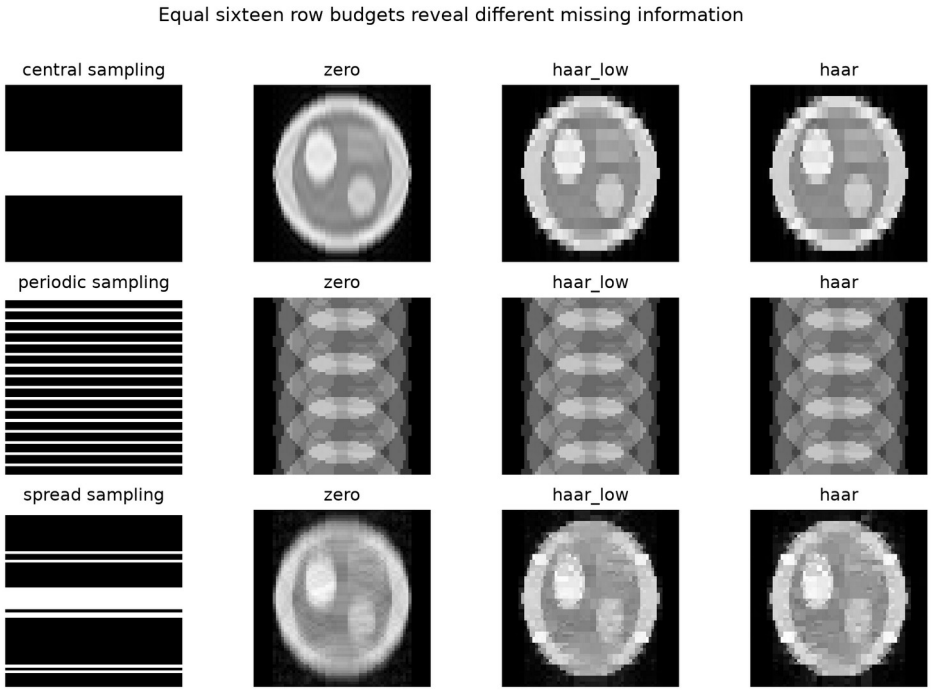


Figure 2. Executed single coil Fourier reconstructions for the first generated object. Rows show central, periodic and spread masks with equal sixteen line budgets; columns show each mask, zero filling, Haar regularization 0.0008 and Haar regularization 0.003. Display ranges are shared.

5 Equal budgets produce unequal recovery

The executed summary in Table 2 demonstrates that global error and bar recovery are related but different outcomes. With central sampling, zero filling gives mean magnitude error of 0.07167 and a mean contrast recovery ratio of 0.035. The standard Haar reconstruction changes the global error to 0.07433 and the recovery ratio to 0.009. These values describe an inverse problem that retains low frequency information but excludes much of the alternating structure. A reconstruction may improve boundary or interior estimates while leaving the specific local signal poorly recovered. The result cannot be attributed to unequal line counts because every acquisition retains exactly sixteen complete rows.

Table 2. Executed mean magnitude RMSE and fine structure recovery ratio. Each row summarizes forty eight objects. Ratios are unitless and error is in arbitrary magnitude units.

Mask	Method	RMSE	Contrast ratio
central	zero	0.07167	0.035
central	haar_low	0.07092	0.007
central	haar	0.07433	0.009
periodic	zero	0.22808	0.142
periodic	haar_low	0.22808	0.142
periodic	haar	0.22808	0.142
spread	zero	0.09314	0.223
spread	haar_low	0.07680	0.253
spread	haar	0.07788	0.290

Periodic sampling produces a different failure geometry. Its measured lines extend across the Fourier grid, but their regular spacing induces coherent spatial ambiguity. The zero filled image therefore contains overlap rather than simply a blurred version of the object. The sparse prior can redistribute image content in response to that overlap, yet its preference for transform sparsity is not an oracle for the generating image. The mean standard Haar error is 0.22808 and its contrast ratio is 0.142. The comparison illustrates why the presence of some high frequencies is not itself sufficient: their arrangement and relation to the object determine what can be separated.

The spread mask preserves a central block and distributes the remaining acquisitions outside it. Under the executed draw, its standard Haar reconstruction has mean magnitude error 0.07788 and mean contrast ratio 0.290. This is evidence about the chosen mask and image family. It is not proof that random spread acquisition always dominates a periodic or low frequency design. A mask aligned with the bar frequency could change the outcome, and a task dominated by broad smooth features could favor another allocation. The appropriate design question is therefore which information the intended target needs under the anticipated object class, rather than which mask has the most appealing general description.

Figure 3 compares the aggregate error and recovery ratio on separate axes. Keeping the scales separate matters because an error decrease and a contrast increase have different meanings and units. The reference

line at a recovery ratio of one marks exact mean local contrast, not an accuracy threshold validated by any reader study. Error bars are not added to imply repeated physical acquisitions; the underlying per case values are instead available in the ledger and summarized through standard deviations in Table 3. This reporting choice favors direct reproducibility over a decorative statistical presentation that could overstate the experiment's population interpretation.

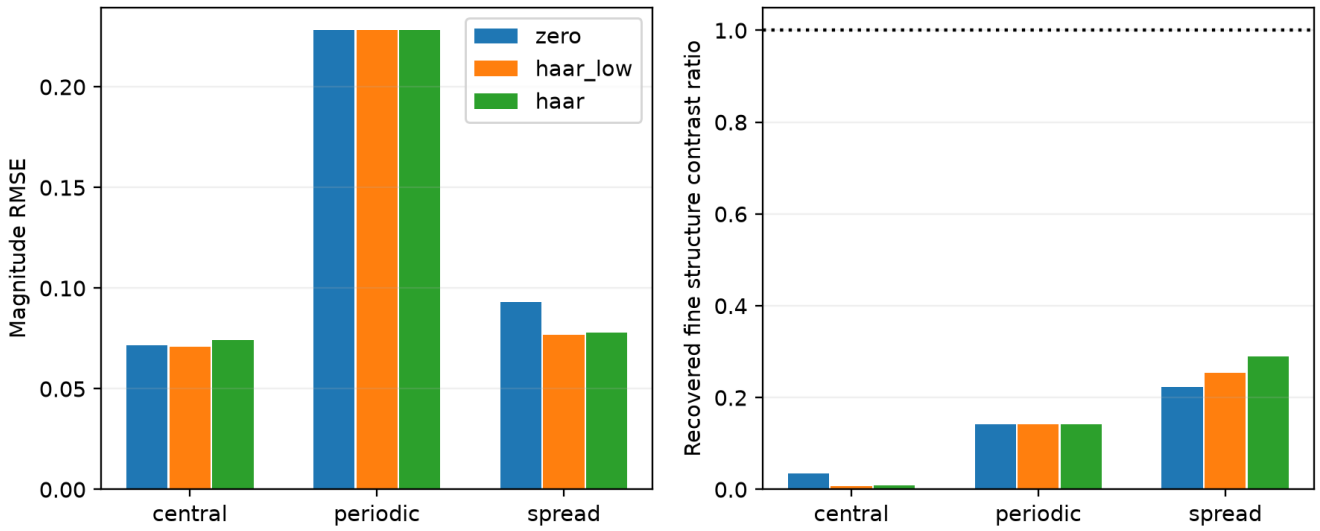


Figure 3. Executed mean magnitude error and fixed bar contrast recovery over forty eight synthetic objects per condition. The dotted contrast reference denotes exact recovery. All estimators retain their acquired noisy Fourier coefficients after final projection.

Changing the shrinkage strength provides a mechanism check. The low penalty and standard penalty use identical measurements, transforms and iterations. Their differences therefore isolate one explicit reconstruction preference. Stronger shrinkage can suppress small transform coefficients that contain either noise or meaningful local detail. The resulting local contrast change need not track the global error change in the same proportion. Table 3 reports the contrast means and standard deviations for these paired settings. This ablation is modest but causally clear within the code: no architecture, training sample, acquisition mask or noise distribution changes when the penalty changes.

Table 3. Local contrast ablation across prescribed penalties. Values are mean recovery ratio and its sample standard deviation over forty eight objects.

Mask	Zero filling	Haar 0.0008	Haar 0.003
central	0.035 ± 0.007	0.007 ± 0.007	0.009 ± 0.007
periodic	0.142 ± 0.020	0.142 ± 0.020	0.142 ± 0.020
spread	0.223 ± 0.016	0.253 ± 0.015	0.290 ± 0.022

The acquired residual remains below 3×10^{-16} across the complete ledger after final projection. That uniform numerical consistency does not erase the differences in Table 2. Instead, it shows that those differences reside in the unmeasured component of the reconstructed image and in magnitude nonlinearities arising from the combined complex field. A report that displayed only the acquired residual would therefore declare these distinct reconstructions similarly successful. The simulation makes that limitation visible

without requiring an adversarial perturbation or a learned model. Even a straightforward deterministic sparse reconstruction can exhibit a mismatch between a consistency check and the local feature a user intends to inspect.

6 Spatial profiles reveal the information tradeoff

The bar region deserves direct spatial inspection because a single contrast ratio can obscure the shape of reconstruction errors. Figure 4 displays matched crops from the standard Haar reconstructions and phase encoding profiles through a fixed horizontal location. The truth profile contains alternating changes superimposed on a varying background. The reconstructions can alter the amplitude, shift apparent local extrema or merge adjacent rows. A mean bar contrast compresses those possibilities into one number. The profile plate therefore serves as a structural diagnostic, allowing the aggregate results to be interpreted in terms of actual pixel arrangements produced by the algorithm.

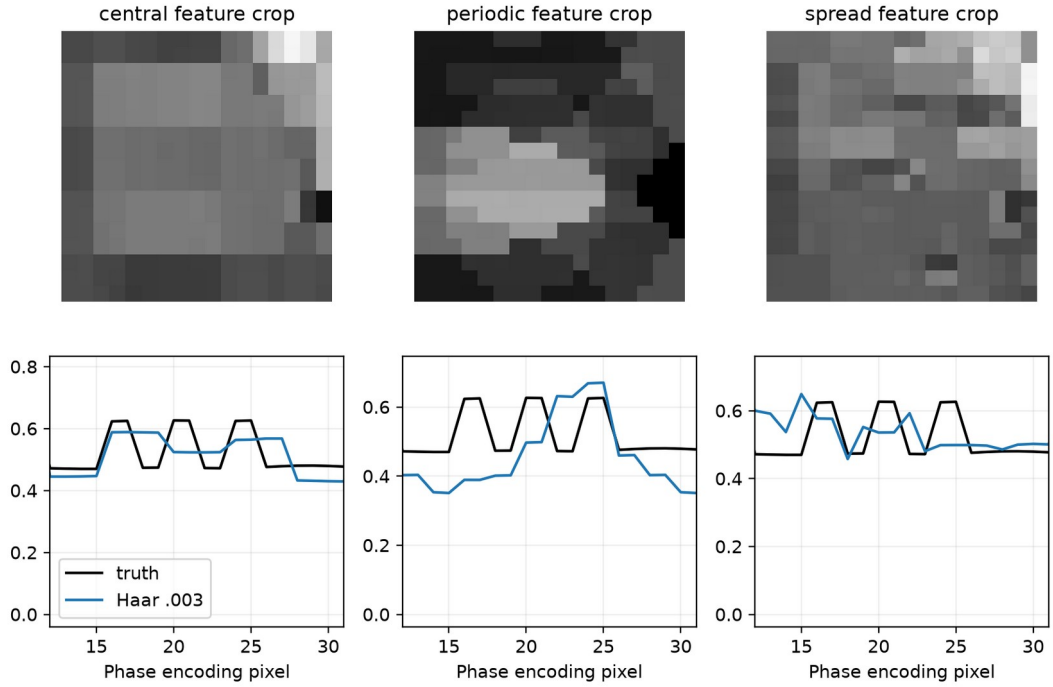


Figure 4. Executed feature crops and fixed column profiles for the standard Haar reconstructions. The generating profile is black. Crops and profiles expose local amplitude loss that a global image error can conceal.

The fixed region definition avoids a favorable selection effect. If the evaluation mask were selected from whichever reconstructed bars appeared strongest, an algorithm could receive credit for a shifted or invented pattern. Here, both positive bar pixels and their interleaved comparison pixels are defined from the generating geometry before reconstruction. The metric penalizes misplaced contrast as well as attenuation. That is appropriate for the stated target, but it also makes the experiment unusually easy in another respect: it assumes perfect knowledge of where the feature should be. Real image interpretation must also decide whether and where an unknown structure exists, which requires a different detection analysis.

Global error includes a large number of uninformative background pixels. Reporting only an object support error would change the weighting of the comparison, but it would still not give the small alternating pattern the same importance as a dedicated regional target. There is no uniquely correct image metric independent of purpose. A reconstruction used for broad anatomical measurement, small vessel depiction or

quantitative parameter fitting may need different evaluations. Our two metrics deliberately expose that dependence rather than attempt to collapse it into a composite score. The bar contrast is a numerical stand in for local detail, not a validated diagnostic endpoint or a claim about any particular disease.

The object phase also matters. Magnitude formation is nonlinear, so errors in complex reconstruction can become amplitude changes in ways that are not captured by a purely real image argument. The smooth phase field is fixed in form but varies slightly across cases, providing limited diversity. We do not estimate phase separately, and no phase correction stage is applied. A broader investigation could vary phase curvature or include coil sensitivities to examine whether additional encoding resolves some of the observed ambiguity. Such an extension would alter the operator and should be evaluated as a new experiment rather than assumed to preserve the same numerical ranking.

Recent reconstruction work continues to explore the relationship between generative priors and measured information. Fast controllable diffusion models study accelerated MRI with diffusion based reconstruction [15], and high frequency conditional diffusion develops another frequency aware approach [16]. Regularized three dimensional diffusion models address volumetric reconstruction [17], while FilterDiff explicitly examines frequency domain diffusion [18]. These studies motivate close attention to where priors enter the reconstruction. Our findings do not contradict their reported benchmark results. They specify a general evaluation issue: a prior's ability to generate plausible detail and a measurement operator's ability to identify that detail are distinct properties requiring distinct evidence.

7 A constructive ambiguity survives perfect acquired agreement

The negative control starts with a rectangular complex image pattern, transforms it to Fourier space, and sets every coefficient selected by the central mask to zero. The inverse transform yields a nonzero image perturbation entirely supported in the unmeasured Fourier subspace. Its maximum absolute amplitude is then scaled to 0.12. Adding that perturbation to the generating image produces an alternative complex image with the same central acquired data apart from numerical rounding. This operation is not an attack on the reconstruction code. It is an explicit construction of the operator's null space, independent of which sparse prior or optimizer is used.

$$h=\alpha F^{-1}[(1-M)Fb], MFh=0 \tag{12}$$

Table 4. Constructive null space control and numerical consistency checks. These values characterize the specified central mask perturbation rather than a patient or clinical feature.

Quantity	Executed value	Meaning
Perturbation maximum	0.12	Prescribed amplitude scale
Image RMS amplitude	0.02203	Nonzero compatible image change
Relative acquired change	5.79×10^{-18}	Essentially zero measured change
Largest estimator residual	2.90×10^{-16}	Exact final acquired replacement

The executed perturbation has image root mean squared amplitude 0.02203, while its measured data difference relative to the complete truth transform is 6×10^{-18}. Table 4 records these values with the construction parameters. The alternative magnitude image visibly differs from the generating magnitude in Figure 5. Consequently, exact agreement on acquired coefficients cannot uniquely determine all image values

for this operator. A prior may favor one member of the compatible set, and an anticipated object family may make that preference useful. The consistency check alone cannot establish that the favored member is the generating object.

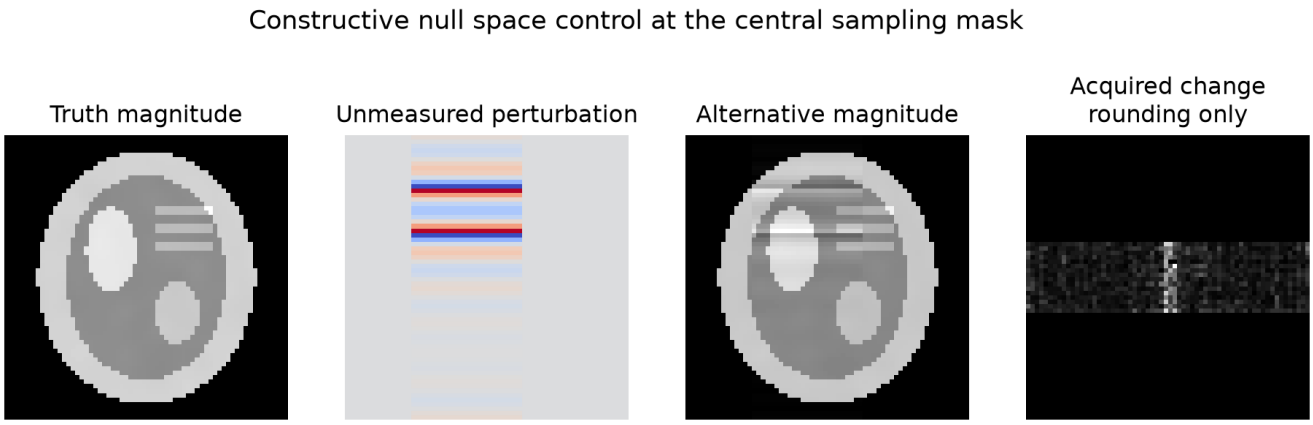


Figure 5. Constructive central mask null space control. A nonzero unmeasured Fourier perturbation changes the compatible magnitude image while its acquired coefficient difference remains at numerical rounding level. The perturbation is an algebraic control rather than an anatomical alternative.

The null space example clarifies the logical status of the empirical findings. The forty eight object comparison estimates how the chosen reconstructions behave within one numerical family. The constructive example proves an algebraic fact about the masked Fourier operator used in that family. It does not prove that every reconstruction must fail on every realistic MRI image, nor that all missing features are unrecoverable under every additional assumption. Multi coil information, stronger support constraints, repeated acquisitions or an accurately specified low dimensional object model can reduce ambiguity. Those are additional information sources. Their value should be demonstrated rather than silently attributed to a small measured residual.

The perturbation is also a warning against interpreting image realism as a substitute for identifiability. Several compatible images can appear reasonable to a viewer, especially when differences are localized or embedded in texture. Conversely, a mathematically compatible image may be anatomically implausible and easily rejected by an effective prior. Our rectangular construction is not offered as a realistic anatomical alternative. Its purpose is narrower: it separates the statement that an image agrees with measured data from the statement that the measured data determine that image. Learned reconstruction evaluations can benefit from retaining this distinction even when the constructed examples themselves are deliberately artificial.

8 Relation to learning and self supervision

No learning occurs in this study. The regularization parameters, Haar transform and mask definitions are prescribed. This is important because references to contemporary deep reconstruction can otherwise make a small numerical experiment sound like an assessment of a trained state of the art model. Self supervised physics guided reconstruction methods use acquired data partitions to train without fully sampled reference images [19]. That strategy changes how model parameters are obtained, but it does not by itself turn a held out frequency subset into a complete test of every spatial feature. The composition of the held out information and the intended downstream target still matter.

ESPIRiT shows how calibration information can define an effective parallel imaging encoding model [20], and deep ESPIRiT reconstruction integrates learned processing with such structure [21]. Those developments

also underline a limitation of our single coil experiment. A real multi coil system can encode spatial information redundantly, so a coefficient missing from one channel need not correspond to the same ambiguity as in our model. We therefore do not infer a clinical sampling recommendation from the spread mask's numerical behavior. The result is a controlled illustration of why comparisons should state the complete operator, rather than summarize it only through an undersampling factor.

Instability analyses of learned image reconstruction identify the possibility of sensitivity to perturbations or missing structures [22]. Our construction examines a different, elementary source of uncertainty: nonuniqueness under incomplete encoding. It would persist even with numerically stable optimization. This distinction matters when selecting a remedy. Better optimization can reduce an algorithmic error, and robustness training can change sensitivity within a model class, but neither automatically acquires information absent from the measurements. A study claiming improved feature preservation should specify whether the improvement comes from more informative acquisition, a more appropriate prior, improved numerical solution or a combination of those mechanisms.

Generative approaches offer expressive image distributions and increasingly flexible inverse solvers. Diffusion posterior sampling explicitly connects sampling with noisy inverse problems [23], while diffusion null space methods emphasize the measured and unmeasured components of an inverse problem [24]. Score based generative modeling provides a broader framework for these constructions [25]. Their relevance here is conceptual rather than computational. A generated sample can express uncertainty or provide plausible completions, but a single sample's visual quality is not proof of correct local detail. Our deterministic comparison provides a simpler setting in which that distinction can be checked without stochastic generation or pretrained model selection.

9 Limitations and implications for acquisition studies

The experiment uses a deliberately small, repeated object family. Its bars have one orientation, one nominal thickness and a restricted amplitude range. A frequency pattern that benefits this target may disadvantage another. The spread mask is drawn only once, so the results cannot estimate the distribution of performance over randomized acquisitions. The image contains no anatomical variation, pathology, motion, chemical shift, relaxation dynamics or receiver sensitivity differences. These omissions are design choices for interpretability, but they substantially limit external validity. The numerical evidence supports statements about the specified Fourier proxy and reconstruction procedure, not about clinical image quality or patient outcomes.

The chosen sparse basis is another material limitation. Haar coefficients efficiently represent some piecewise constant components but are not a universal representation of MRI anatomy or alternating texture. A different transform, total variation penalty or learned prior could change both global error and local recovery. The final data projection also imposes exact agreement with noisy measurements, which may be suboptimal at higher noise. We retain it to make the consistency argument clear, not because it is asserted to be the best denoising strategy. The comparison should therefore be read as a controlled investigation of interacting choices rather than an optimization contest.

A useful next acquisition study would vary feature orientation and frequency before choosing a reconstruction parameter. That design would reveal whether a mask's favorable performance reflects general diversity or accidental alignment with one local target. It could also include multiple independent mask draws, paired noise levels and a fully sampled noisy reference to separate measurement noise from omission. These are concrete extensions of the present numerical logic. They would require new executed evidence

before supporting broader claims. We do not infer their outcomes or present them as if they had already been measured in the archived simulation.

The practical reporting implication is straightforward. An acceleration comparison should give the encoding operator, the reconstruction preference and a target specific evaluation alongside a global image score. Acquired data residuals remain necessary checks for implementation errors, but their interpretation should stop at the measured subspace. The present experiment demonstrates why: all methods can have negligible residuals while recovering different local contrast, and a nonzero perturbation can leave acquired data unchanged. Equal measurement budgets are therefore a controlled resource constraint, not a guarantee of equal information. Within this synthetic study, fine structure preservation depends on which measurements are retained and how the reconstruction completes what remains unobserved.

10 The unchanged periodic solution is an informative negative result

One result warrants more attention than a conventional ranking would give it. For the periodic mask, both Haar penalties return the same final magnitude error and contrast ratio as zero filling to the precision reported in the ledger. The reconstruction procedure has not discovered useful missing information in this condition. The periodic sampling operator creates a repeated alias pattern aligned with the dyadic structure of the Haar representation. Starting from zero filling, shrinkage can preserve that repeated structure, and the final acquired coefficient replacement restores the observed component. The output can consequently remain in the same restricted Fourier subspace instead of synthesizing a discriminating completion.

This is an algorithm and acquisition interaction, rather than evidence that sparsity has no value in MRI. The same code changes the spread reconstruction substantially. Its behavior depends on whether the update can leave the symmetry inherited from the initial estimate. A learned model, another transform, an asymmetric initialization or a different phase encoding mask could break that symmetry. Such changes would need to be evaluated explicitly. The executed negative result is useful precisely because it blocks an overly broad narrative in which iterative processing must improve every undersampled input. Eighty iterations and an established optimization recurrence do not guarantee that the chosen representation contributes new information for every operator.

The central mask presents a complementary limitation. The standard penalty slightly increases global error relative to zero filling while leaving recovered alternating contrast very small. It is therefore inappropriate to present the study as a general success of stronger regularization. At the lower penalty the global error decreases slightly, yet local contrast remains close to zero. The spread mask supports a different combination: both sparse settings improve global error compared with its zero filled baseline, and the standard penalty increases recovered bar contrast relative to the lower penalty. These outcomes show that the effect of shrinkage cannot be summarized as uniform smoothing of all detail. It depends on how the missing data and transform preferences interact.

The fine structure is not fully recovered even in the most favorable reported condition. A mean ratio around 0.29 means that much of the generating alternating contrast is still missing. Calling that condition successful without the ratio would overstate the evidence. The comparison establishes a relative preservation difference at a common measurement budget, not faithful reconstruction of the complete object. It also explains why a favorable looking spread image should be interpreted cautiously. Broad components can appear convincing while the deliberately measured fine pattern remains substantially attenuated. The image plate and numerical target together provide a more accurate account than either one alone.

The spread design is also not the lowest global error design across every method. Central low penalty reconstruction gives a smaller global magnitude error than the standard spread reconstruction in this particular experiment, while its bar ratio is much smaller. That crossing is the central task tradeoff rather than an inconvenient exception. If the application cared only about the aggregate magnitude score, the central allocation could be preferred. If preserving the alternating difference carried greater value, the spread allocation would provide more of that target information. The choice cannot be resolved by declaring one metric universally superior; it requires a stated use and an evaluation aligned with that use.

The distinction between conditioning and completion also connects to multi domain diffusion prior guidance [26]. That work combines image, latent and frequency information within a much richer learned method. Our study neither reproduces its experiments nor estimates its performance. It supplies a simple baseline question that remains useful when such architectures are evaluated: which improvements are supported directly by acquired coefficients, and which depend on the learned distribution used to complete missing information. Answering that question requires local targets or carefully designed controls in addition to mean image quality, particularly when a reconstruction will be interpreted as evidence for a small spatial structure.

A further methodological consequence concerns how negative controls are selected. A random image perturbation usually changes both acquired and unacquired data, so a reconstruction can reject it using a nonzero residual. The projected perturbation in Figure 5 removes that easy route by construction. Its compatibility is an algebraic property, not an accidental result from a small sample. That makes it suitable for testing the interpretation of consistency, although not for measuring the prevalence of realistic ambiguous anatomy. A strong evaluation can include both types of evidence: representative simulations for typical behavior and constructive controls for the limits of a claimed diagnostic criterion.

The archived evidence permits these interpretations to be checked without relying on the prose. Each of the 432 ledger entries records its object index, mask, method, exact and reconstructed contrast, recovery ratio, magnitude error and acquired residual. The script regenerates the same object sequence and fixed spread mask from seed 81125. The figures derive from those operations, with no external data or manually corrected result. This traceability is especially helpful for the unchanged periodic output, because a reader can inspect whether the equality is a reporting omission or an actual numerical outcome. In this experiment it is an actual outcome, and it constrains the scientific conclusion accordingly.

References

[1] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, Peter Boesiger. SENSE: Sensitivity encoding for fast MRI. Magnetic Resonance in Medicine, 1999. doi:10.1002/(sici)1522-2594(199911)42:5<952::aid-mrm16>3.0.co;2-s.

[2] Daniel K. Sodickson, Warren J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. Magnetic Resonance in Medicine, 1997. doi:10.1002/mrm.1910380414.

[3] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, Axel Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magnetic Resonance in Medicine, 2002. doi:10.1002/mrm.10171.

[4] Michael Lustig, David Donoho, John M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magnetic Resonance in Medicine, 2007. doi:10.1002/mrm.21391.

[5] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, Florian Knoll. Learning a variational network for reconstruction of accelerated MRI data. Magnetic Resonance in Medicine, 2017. doi:10.1002/mrm.26977.

[6] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony N. Price, Daniel Rueckert. A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction. IEEE Transactions on Medical Imaging, 2018. doi:10.1109/tmi.2017.2760978.

[7] Hemant Kumar Aggarwal, Merry P. Mani, Mathews Jacob. MoDL: Model Based Deep Learning Architecture for Inverse Problems. arXiv preprint arXiv:1712.02862v1, 2017. https://arxiv.org/abs/1712.02862v1
[8] Jure Zbontar et al.. fastMRI: An Open Dataset and Benchmarks for Accelerated MRI. arXiv preprint arXiv:1811.08839v1, 2018. https://arxiv.org/abs/1811.08839v1
[9] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C. Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, Patricia Johnson. End-to-End Variational Networks for Accelerated MRI Reconstruction. Lecture Notes in Computer Science, 2020. doi:10.1007/978-3-030-59713-9_7.
[10] Amir Beck, Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. SIAM Journal on Imaging Sciences, 2009. doi:10.1137/080716542.
[11] Leonid I. Rudin, Stanley Osher, Emad Fatemi. Nonlinear total variation based noise removal algorithms. Physica D: Nonlinear Phenomena, 1992. doi:10.1016/0167-2789(92)90242-f.
[12] Singanallur V. Venkatakrishnan, Charles A. Bouman, Brendt Wohlberg. Plug-and-Play priors for model based reconstruction. 2013 IEEE Global Conference on Signal and Information Processing, 2013. doi:10.1109/globalsip.2013.6737048.
[13] Yaniv Romano, Michael Elad, Peyman Milanfar. The Little Engine That Could: Regularization by Denoising (RED). SIAM Journal on Imaging Sciences, 2017. doi:10.1137/16m1102884.
[14] Victor Lempitsky, Andrea Vedaldi, Dmitry Ulyanov. Deep Image Prior. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018. doi:10.1109/cvpr.2018.00984.
[15] Wei Jiang, Zhuang Xiong, Feng Liu, Nan Ye, Hongfu Sun. Fast Controllable Diffusion Models for Undersampled MRI Reconstruction. 2024 IEEE International Symposium on Biomedical Imaging (ISBI), 2024. doi:10.1109/isbi56570.2024.10635891.
[16] Ziwei Zhu, Lijun Ouyang, Dongping Xiong, Xiaozhi Zhang. High-Frequency Conditional Diffusion Probabilistic Model for Under-sampled MRI Reconstruction. 2024 International Conference on Image Processing, Computer Vision and Machine Learning (ICICML), 2024. doi:10.1109/icicml63543.2024.10958164.
[17] Arya Bangun, Zhuo Cao, Alessio Quercia, Hanno Scharr, Elisabeth Pfaehler. MRI Reconstruction with Regularized 3D Diffusion Model (R3DM). 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 2025. doi:10.1109/wacv61041.2025.00078.
[18] Tao Song, Fang Nie, Yi Guo, Feng Xu, Shaoting Zhang. FilterDiff: Noise-Free Frequency-Domain Diffusion Models for Accelerated MRI Reconstruction. Lecture Notes in Computer Science, 2025. doi:10.1007/978-3-032-05325-1_20.
[19] Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, Jutta Ellermann, Kâmil Uğurbil, Mehmet Akçakaya. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. Magnetic Resonance in Medicine, 2020. doi:10.1002/mrm.28378.
[20] Martin Uecker, Peng Lai, Mark J. Murphy, Patrick Virtue, Michael Elad, John M. Pauly, Shreyas S. Vasanawala, Michael Lustig. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA. Magnetic Resonance in Medicine, 2013. doi:10.1002/mrm.24751.
[21] Christopher M. Sandino, Peng Lai, Shreyas S. Vasanawala, Joseph Y. Cheng. Accelerating cardiac cine MRI using a deep learning-based ESPIRiT reconstruction. Magnetic Resonance in Medicine, 2020. doi:10.1002/mrm.28420.
[22] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, Anders C. Hansen. On instabilities of deep learning in image reconstruction and the potential costs of AI. Proceedings of the National Academy of Sciences, 2020. doi:10.1073/pnas.1907377117.
[23] Hyungjin Chung, Jeongsol Kim, Michael T. McCann, Marc L. Klasky, Jong Chul Ye. Diffusion Posterior Sampling for General Noisy Inverse Problems. arXiv preprint arXiv:2209.14687v3, 2023. https://arxiv.org/abs/2209.14687v3
[24] Yinhuai Wang, Jiwen Yu, Jian Zhang. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. arXiv preprint arXiv:2212.00490v2, 2022. https://arxiv.org/abs/2212.00490v2
[25] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. arXiv preprint arXiv:2011.13456v2, 2021. https://arxiv.org/abs/2011.13456v2
[26] Lingtong Zhang, Mengdie Song, Xiaohan Hao, Huayu Mai, Bensheng Qiu. MDPG: Multi-domain Diffusion Prior Guidance for MRI Reconstruction. arXiv preprint arXiv:2506.23701v1, 2025. https://arxiv.org/abs/2506.23701v1