Artificial Intelligence Methods for MRI Reconstruction: From Image-Domain Denoising to Physics-Informed, Clinically Robust Imaging

B

May 2026

This document is a personal technical summary and interpretation of selected published work. All figures and graphical illustrations intended for this document have been generated separately for technical explanation. They should not reproduce or knowingly imitate copyrighted figures from published literature, commercial documentation, or other protected sources. No third-party copyrighted images are used or claimed.

Contents

1 Introduction: Why Artificial Intelligence Entered MRI Reconstruction

Magnetic resonance imaging (MRI; MRT in German usage) is clinically valuable because it provides excellent soft-tissue contrast without ionizing radiation. Its central technical limitation is acquisition time. MRI scanners do not directly acquire images; rather, they measure spatial-frequency data in k-space, and the number of acquired k-space samples is closely linked to scan duration. Accelerated MRI therefore seeks to reduce acquisition time by undersampling k-space. However, undersampling renders the reconstruction problem ill-posed and can produce aliasing artifacts unless the missing information is recovered through suitable reconstruction constraints or learned priors [1–3].

Before deep learning became prominent, MRI acceleration already relied heavily on parallel imaging and compressed sensing. Parallel imaging exploits the spatially varying sensitivity profiles of multiple receiver coils, whereas compressed sensing reconstructs undersampled data by combining transform-domain sparsity with suitably incoherent sampling. In more recent reconstruction research, deep learning has increasingly been integrated into the inverse problem itself: it can supply learned priors and regularizers, replace or augment iterative reconstruction updates, and, in some approaches, even optimize the k-space sampling strategy jointly with reconstruction [4–7].

The central thesis of this report is that the strongest AI methods for MRI reconstruction are not unconstrained black boxes. The field has increasingly moved from purely image-domain artifact removal toward physics-informed learned reconstruction, in which acquisition models and measured data remain structurally embedded in the reconstruction process. Key developments include data-consistency mechanisms, explicit use of multi-coil information, k-space and dual-domain operations, unrolled optimization schemes, self-supervised training without fully sampled reference data, adaptive or learned sampling strategies, and increasingly rigorous robustness assessment through cross-domain testing, radiologist evaluation, and analysis of hallucination risks [7–12].

Figure 1 summarizes this overall progression, from acquisition constraints and undersampled k-space to learned reconstruction, robustness assessment, and clinical deployment.

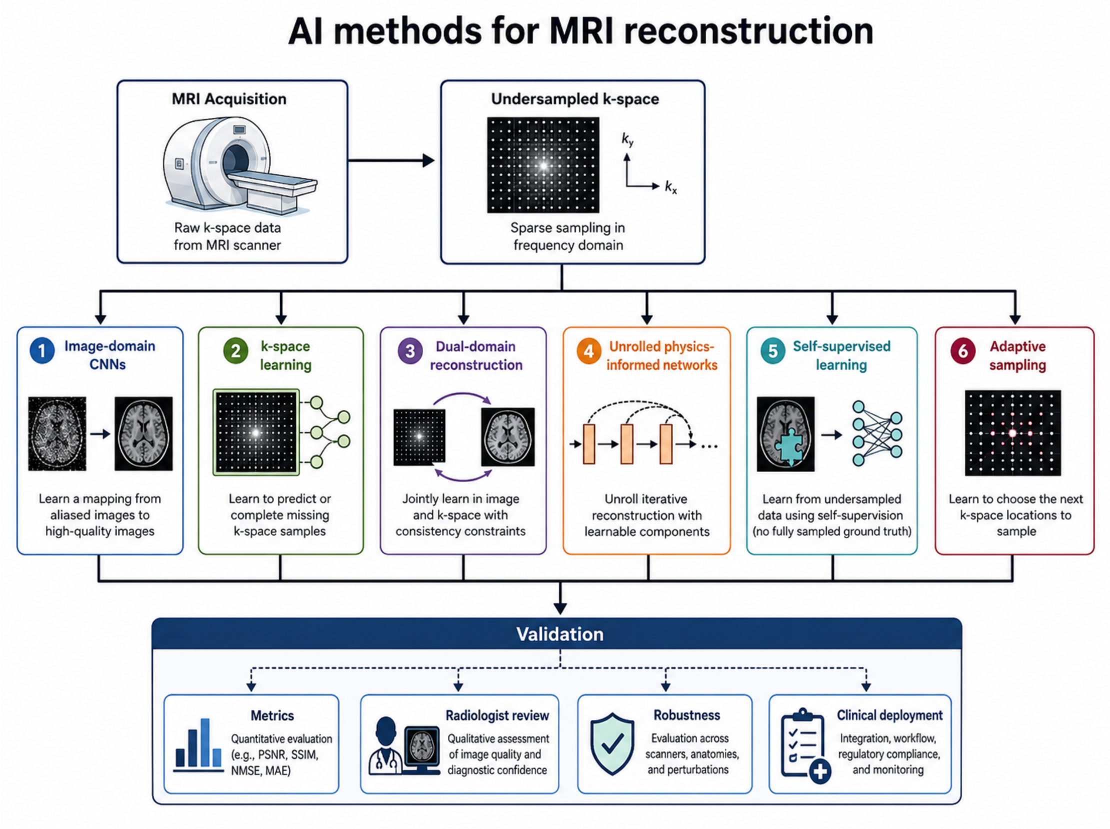

Figure 1: Conceptual landscape of AI methods in MRI reconstruction. The figure should show the relationship between acquisition, k-space undersampling, classical reconstruction, learned reconstruction, robustness evaluation, and clinical deployment.

2 MRI Reconstruction as a Physics-Constrained Inverse Problem

2.1 K-space, Fourier Encoding, and Data Fidelity

In a simplified setting, MRI reconstruction estimates an image x from measured k-space data y. With undersampling, the measurement equation can be written schematically as

$$y = MFx + \varepsilon, \tag{1}$$

where F denotes Fourier encoding, M is the sampling mask, and ε represents noise and imperfections. When M removes part of k-space, the reconstruction problem becomes underdetermined and, more generally, ill-posed: the acquired measurements no longer uniquely determine the image without additional constraints or prior information. Modern reconstruction methods therefore should not merely generate visually plausible images; they must also preserve consistency with the actually measured k-space samples [1, 8, 10].

Figure 2 illustrates this inverse-problem structure, including the transition from complete encoding to undersampled k-space, aliased intermediate images, and data-consistent learned reconstruction.

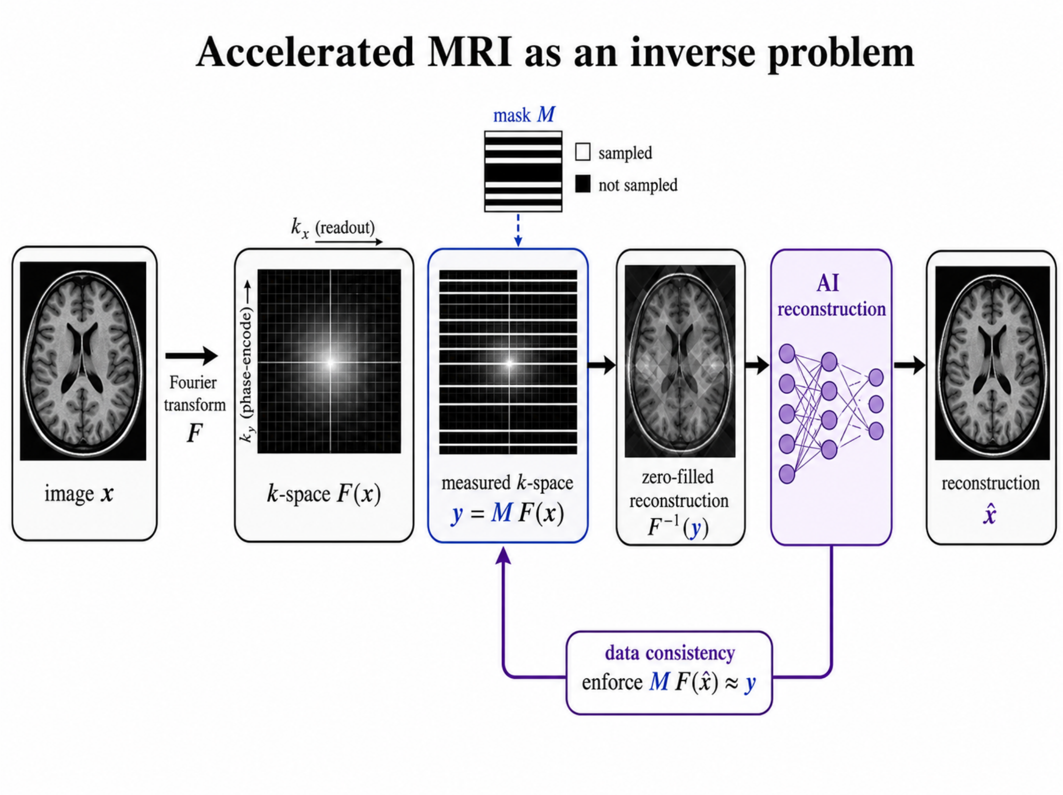

Figure 2: MRI reconstruction as a physics-constrained inverse problem. A suitable schematic should show an anatomical image, Fourier/k-space encoding, undersampling, aliasing in the zero-filled reconstruction, and learned reconstruction constrained by measured k-space.

2.2 Parallel Imaging and Compressed Sensing

Modern clinical MRI typically uses multiple receiver coils. Each coil measures the anatomy through a distinct spatial sensitivity profile, and parallel imaging exploits the resulting redundancy across coils to reconstruct images from undersampled k-space data. In the classical taxonomy, GRAPPA-type methods estimate missing information locally in k-space, whereas SENSE-type methods use explicit coil-sensitivity modeling in the image domain. Deep learning methods such as GrappaNet show that these established parallel-imaging principles can be incorporated directly into neural reconstruction pipelines, for example by embedding a differentiable GRAPPA operation within an end-to-end trainable model [4, 13].

Compressed sensing provides a second foundation for accelerated MRI. It reconstructs images from undersampled data by imposing prior structure, classically sparsity in a suitable transform domain. The relationship between compressed sensing and deep learning should not be cast as a sharp conceptual break: many learned reconstruction methods preserve the same regularized inverse-problem structure, retaining data-consistency terms while replacing hand-designed sparsity priors or proximal operators with trainable regularizers [5, 6, 14].

2.3 Why MRI AI Is Not Ordinary Denoising

Ordinary denoising starts from a corrupted image, whereas MRI reconstruction starts from incomplete physical measurements. This distinction matters clinically: a reconstruction may appear visually sharp yet remain unreliable if it is insufficiently constrained by the acquired k-space data, suppresses small pathologies, or introduces plausible but false structure. For this reason, the assessment of AI-based MRI reconstruction must extend beyond perceptual image quality to include data consistency, stability under perturbations and distribution shifts, recovery of diagnostically relevant fine detail, and explicit analysis of hallucination risks [8, 12, 15, 16].

3 From Classical Reconstruction to Learned Reconstruction

3.1 Classical Methods Remain Conceptually Central

A mature account of AI in MRI should avoid a simplistic "classical methods versus deep learning" narrative. The paper revisiting ℓ_1-wavelet compressed-sensing MRI is especially important in this regard because it shows that, when compressed sensing is refined using tools common in contemporary deep-learning reconstruction—including algorithm unrolling, end-to-end tuning over a training dataset, and modern optimization strategies—it can perform close to a physics-guided deep-learning method while using far fewer tunable parameters [6].

The implication is that part of the performance advantage often attributed broadly to "deep learning" may in some comparisons reflect data-driven parameter tuning, algorithm unrolling, and modern optimization practices rather than the mere replacement of classical reconstruction principles. Deep learning remains powerful, but its most convincing role in accelerated MRI reconstruction is often to extend classical inverse-problem formulations—for example through learned regularizers embedded within physics-based, data-consistent reconstruction schemes—rather than to discard them altogether [6, 8].

3.2 AI as Learned Regularization

A useful way to understand many AI methods for MRI reconstruction is as forms of learned regularization. Classical compressed sensing constrains the ill-posed reconstruction problem through hand-designed priors, such as transform-domain sparsity or total-variation regularization. Learned reconstruction methods replace or augment such handcrafted priors with trainable regularization networks, while explicit data-consistency terms preserve the MRI forward model and enforce agreement with the acquired measurements [5, 6, 8].

Figure 3 contrasts classical compressed-sensing regularization with learned physics-informed reconstruction and highlights their shared reliance on consistency with measured k-space data.

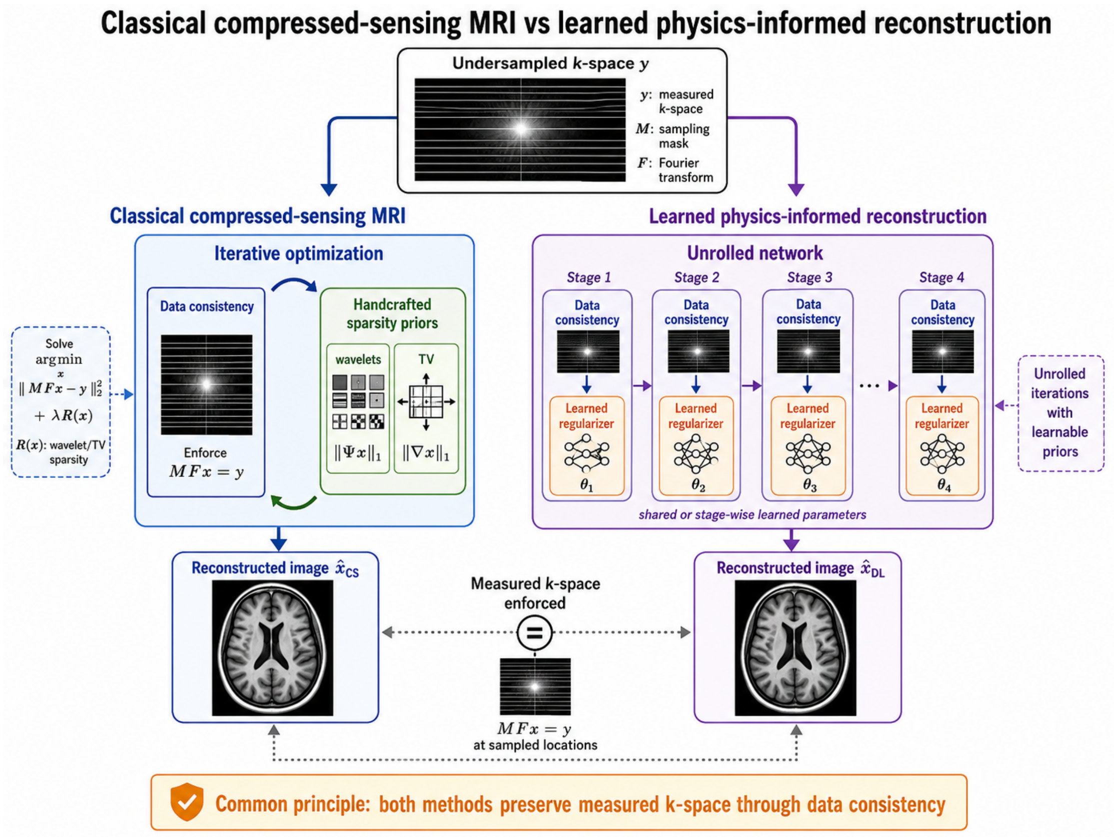

Figure 3: From classical to learned regularization. The figure should contrast a classical compressed-sensing pipeline with a learned unrolled reconstruction pipeline, emphasizing that both use data consistency but differ in the regularizer or update module.

4 Taxonomy of AI Methods in MRI Reconstruction

4.1 Image-Domain CNN De-Aliasing

Early deep-learning MRI reconstruction often treated the task as image-domain artifact removal. A zero-filled or otherwise aliased reconstruction was passed into a convolutional neural network, which learned to map it toward a cleaner target image. This approach was historically important because it demonstrated that learned nonlinear mappings could substantially improve upon aliased baseline reconstructions while enabling fast feed-forward inference [1, 3, 17].

Its limitation is equally important: purely image-domain learning can become weakly connected to the acquired measurements. Without mechanisms that enforce consistency with measured k-space, a network may suppress genuine structures or generate visually plausible but unsupported details. Concerns of this kind helped motivate architectures that more tightly integrate the measurement model, including cascaded reconstructions with explicit data-consistency steps, dual-domain networks that alternate between image and k-space representations, and unrolled reconstruction schemes that embed learned components within iterative inverse-problem solvers [8, 10, 15, 18].

4.2 Data Consistency Layers

Data consistency is one of the defining ideas of MRI-specific AI reconstruction. It constrains the learned reconstruction by enforcing agreement with the acquired k-space measurements, either by preserving sampled values directly or by penalizing deviations from them. The systematic evaluation of iterative deep networks identifies data-consistency layers and expressive regularization networks as core components of robust fast parallel MRI reconstruction [1, 8].

Data consistency appears across several major MRI reconstruction families, including cascaded CNNs, unrolled networks, and dual-domain architectures. DuDoRNet, for example, interleaves image-domain restoration, k-space-domain restoration, and data-consistency operations within recurrent reconstruction blocks. GAN-based frameworks have also incorporated measurement-fidelity constraints, although the exact form of those constraints varies by method [8, 10, 18, 19].

4.3 K-space Learning and Parallel MRI Learning

AI reconstruction becomes particularly MRI-native when it operates directly on k-space measurements and exploits multi-coil information rather than treating reconstruction as generic image enhancement. GrappaNet is a key example. It combines a differentiable GRAPPA-style layer with neural networks, processes complex-valued multi-coil views jointly before final coil combination, and refines the reconstruction in both k-space and image space [4, 13].

Figure 4 depicts this multi-coil reconstruction logic, linking receiver-coil sensitivity information, learned k-space completion, image-domain refinement, and final coil-combined output.

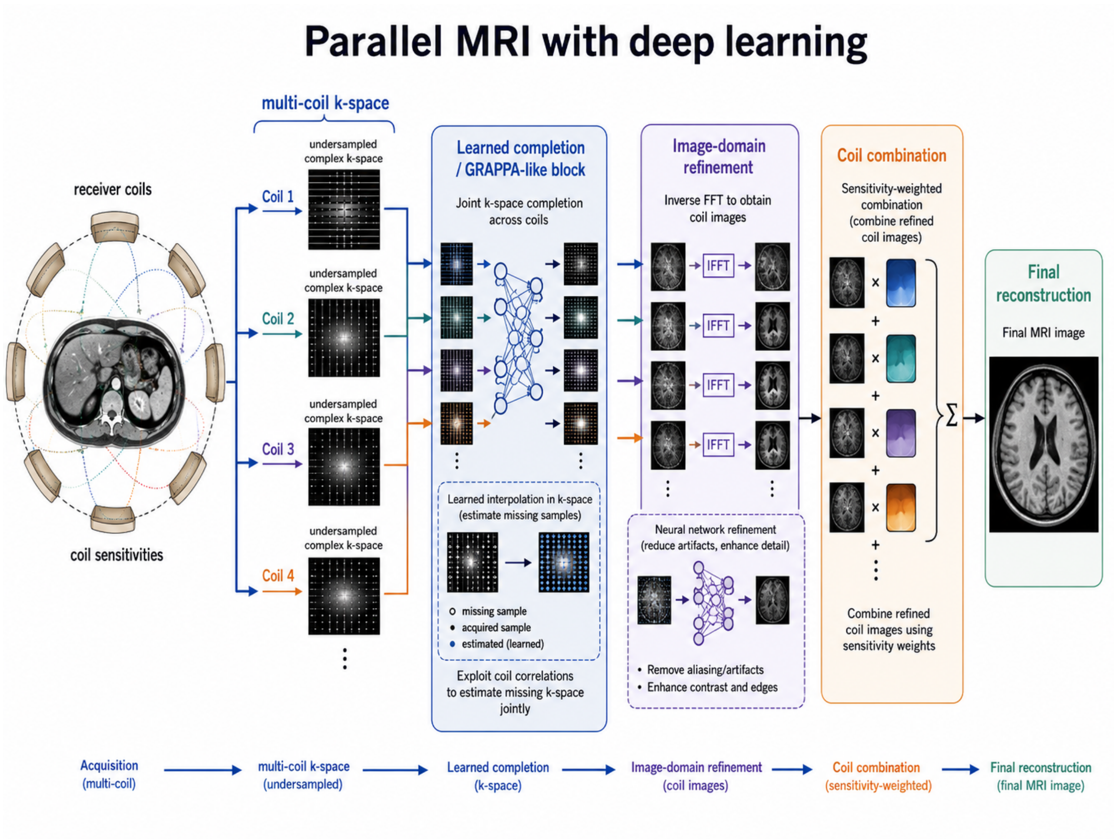

Figure 4: Parallel MRI with deep learning. The figure should show multiple receiver coils, undersampled multi-coil k-space, coil-sensitive reconstruction, a GRAPPA-like or learned k-space completion block, and final coil-combined image.

4.4 Dual-Domain and Cross-Domain Reconstruction

Dual-domain methods exploit complementary information from both domains: spatial and anatomical structure in the image domain, and frequency-domain acquisition information in k-space. DuDoRNet is the clearest example in the literature: each recurrent block alternates between image-domain restoration, Fourier transformation, data consistency, k-space restoration, a second data-consistency step, and inverse Fourier reconstruction [1, 10].

PIC-GAN also reflects the broader trend of integrating adversarial learning with MRI-specific reconstruction structure. It couples parallel imaging with a GAN-based reconstruction framework for accelerated multi-channel MRI [19]. More generally, dual-domain learning is a natural maturation of MRI AI: the network participates in reconstruction across the image and k-space domains that define the modality, rather than merely cleaning the final image [1, 10].

Figure 5 visualizes this recurrent dual-domain pattern, in which image-space and k-space updates are linked by Fourier transforms and repeated data-consistency operations.

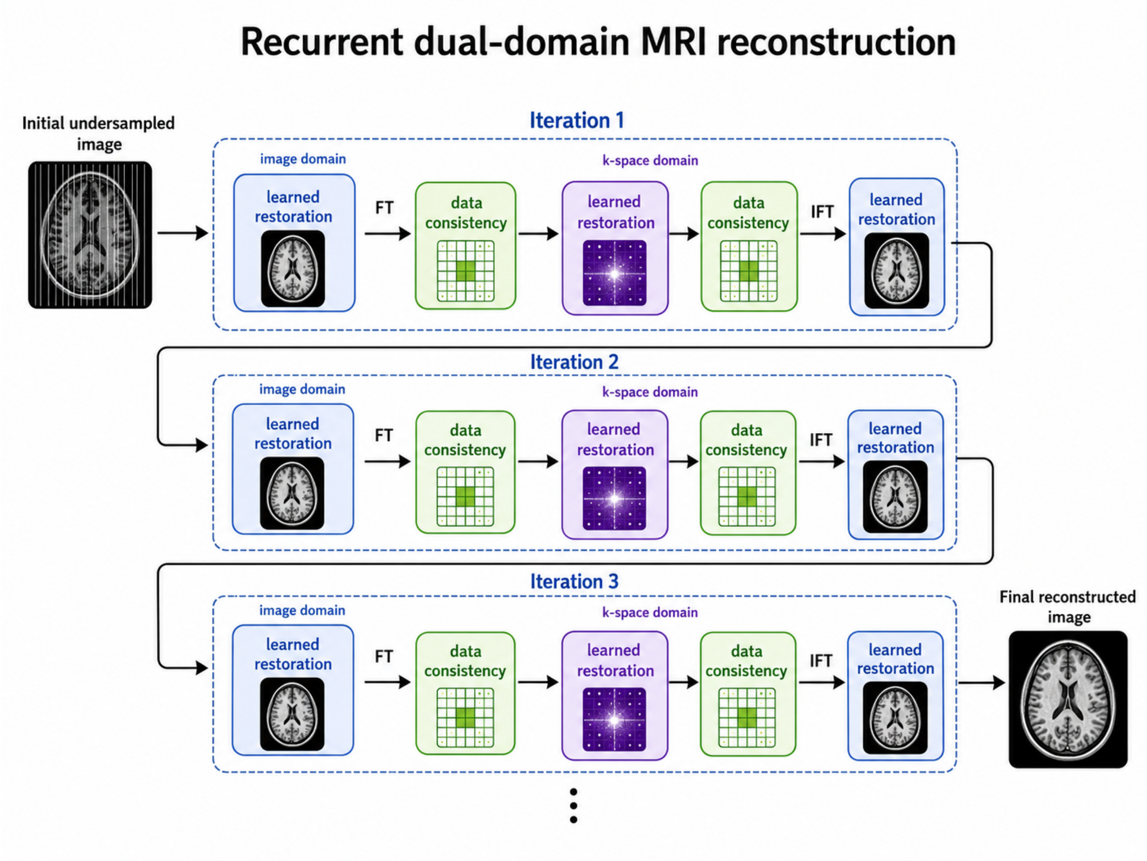

Figure 5: Dual-domain MRI reconstruction. The figure should show recurrent alternation between image-domain restoration and k-space-domain restoration, with Fourier transforms and data-consistency blocks linking the two domains.

4.5 Physics-Informed and Unrolled Networks

Unrolled networks are among the most important AI families for accelerated MRI. They reinterpret iterative reconstruction as a trainable architecture: a finite number of optimization iterations is unfolded into successive network blocks, learned modules provide regularization-like updates, and data-consistency operations preserve the MRI forward model. RecurrentVarNet is a strong example because it performs recurrent refinement for accelerated multi-coil MRI and carries out iterative optimization in k-space [1, 8, 9].

The systematic evaluation paper is a major anchor for this subsection. It compares iterative reconstruction networks under varying data-consistency layers, regularization networks, training sample sizes, anatomies, and acceleration factors. It reports that physics-based reconstruction networks substantially outperform pure post-processing approaches at $R = 4$, while alignment between training and test domains becomes especially important at the higher acceleration factor $R = 8$ [8].

Figure 6 summarizes the unrolling principle by showing a finite sequence of learned regularization steps and physics-based data-consistency updates.

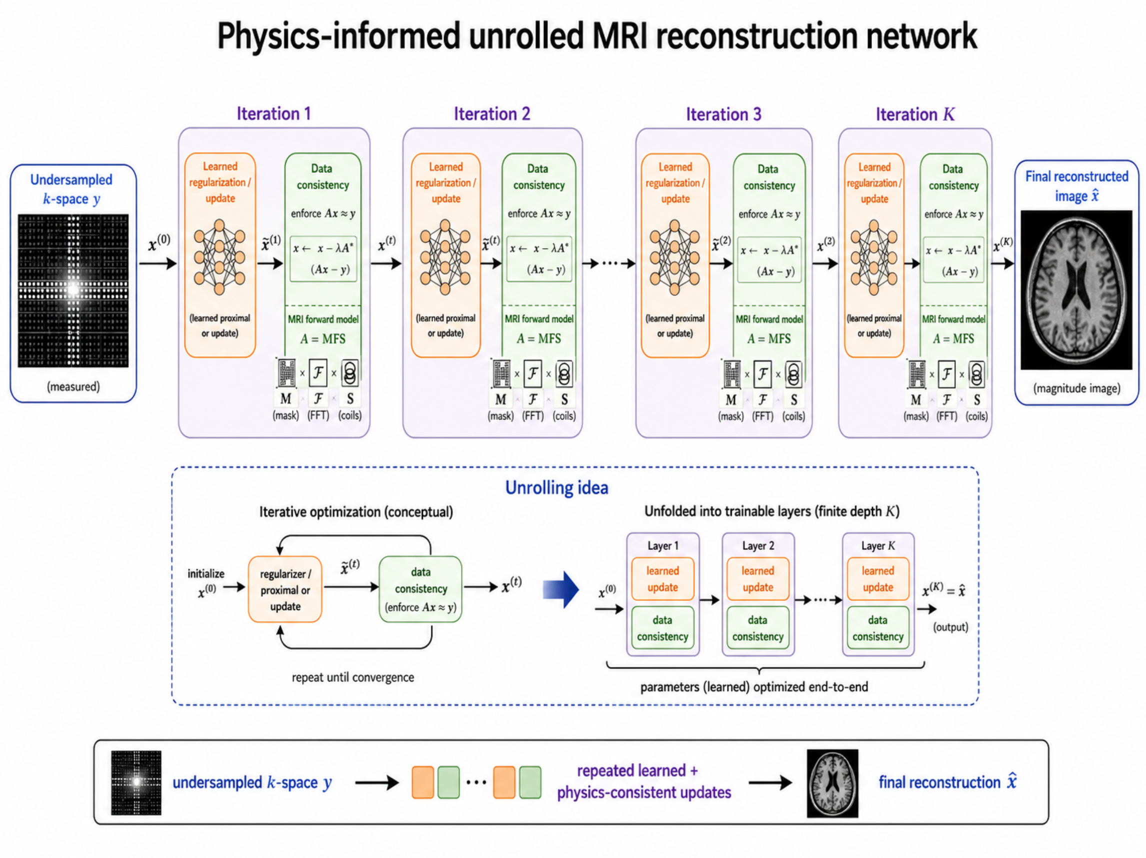

Figure 6: Physics-informed unrolled reconstruction. The figure should show repeated learned regularization/update blocks followed by data-consistency blocks, resembling an iterative reconstruction algorithm unfolded into a neural network.

4.6 GANs and Perceptual Reconstruction

GAN-based reconstruction and perceptual-loss strategies can favor visually sharper or more detail-preserving images, but they require careful treatment in diagnostic MRI. PIC-GAN is relevant because it couples adversarial learning with parallel imaging and combines adversarial loss with pixel-wise fidelity losses in both image and frequency domains [19]. The broader caution is that perceptual plausibility must not be confused with truthfulness to the measured data. The instability, robustness, and fastMRI challenge literature show why visually convincing reconstructions must be evaluated for perturbation sensitivity, preservation of clinically relevant structures, hallucination risk, and robustness under domain shift [12, 15, 16].

4.7 Self-Supervised, Unsupervised, and No-Ground-Truth Learning

Fully sampled reference data are difficult or impractical to obtain in several MRI settings. This motivates self-supervised and unsupervised reconstruction methods, especially approaches that learn from undersampled acquisitions without requiring fully sampled k-space targets. MRI-specific work should remain central in this section, including reconstruction methods designed explicitly for the absence of fully sampled reference data and unsupervised MRI applications such as ASL reconstruction. More general denoising work on learning without clean paired targets may serve as conceptual background, but should not define the MRI reconstruction argument [20, 21].

The key point is that general self-supervised denoising ideas do not automatically solve MRI reconstruction. MRI methods must still address k-space sampling, data consistency, complex-valued data, and multi-coil physics.

4.8 Multi-Contrast and Anatomical-Prior-Guided MRI Reconstruction

Multi-contrast or multi-protocol MRI examinations often acquire several related image contrasts within the same session. AI reconstruction can exploit this structure by using a shorter or already available contrast as prior information for reconstructing another. DuDoRNet uses fully sampled T1-weighted information as a deep prior to guide reconstruction of longer protocols, embedding that prior in both image and k-space domains [10]. Deep-learning-based multi-modal fusion for fast MR reconstruction similarly combines fully sampled T1-weighted information with undersampled T2-weighted data to support reconstruction of the target T2-weighted image [22].

This is one of the most MRI-specific AI strategies discussed here. It is promising because related MR contrasts share anatomical information while providing complementary tissue contrast, making one sequence potentially useful as prior information for reconstructing another [10, 22]. At the same time, this strategy raises an important caution: if pathology is visible only in one sequence, or manifests differently across contrasts, a prior-guided reconstruction could in principle bias the target reconstruction toward structures suggested by the prior image. This concern should be interpreted in light of the broader instability and robustness literature, which emphasizes sensitivity to small structural changes, preservation of diagnostically relevant detail, and the need to test learned reconstruction methods under clinically meaningful distribution shifts [15, 16].

Figure 7 illustrates the central idea of prior-guided reconstruction: a fully sampled anatomical contrast can assist reconstruction of an undersampled target contrast, while also motivating caution about prior-induced bias.

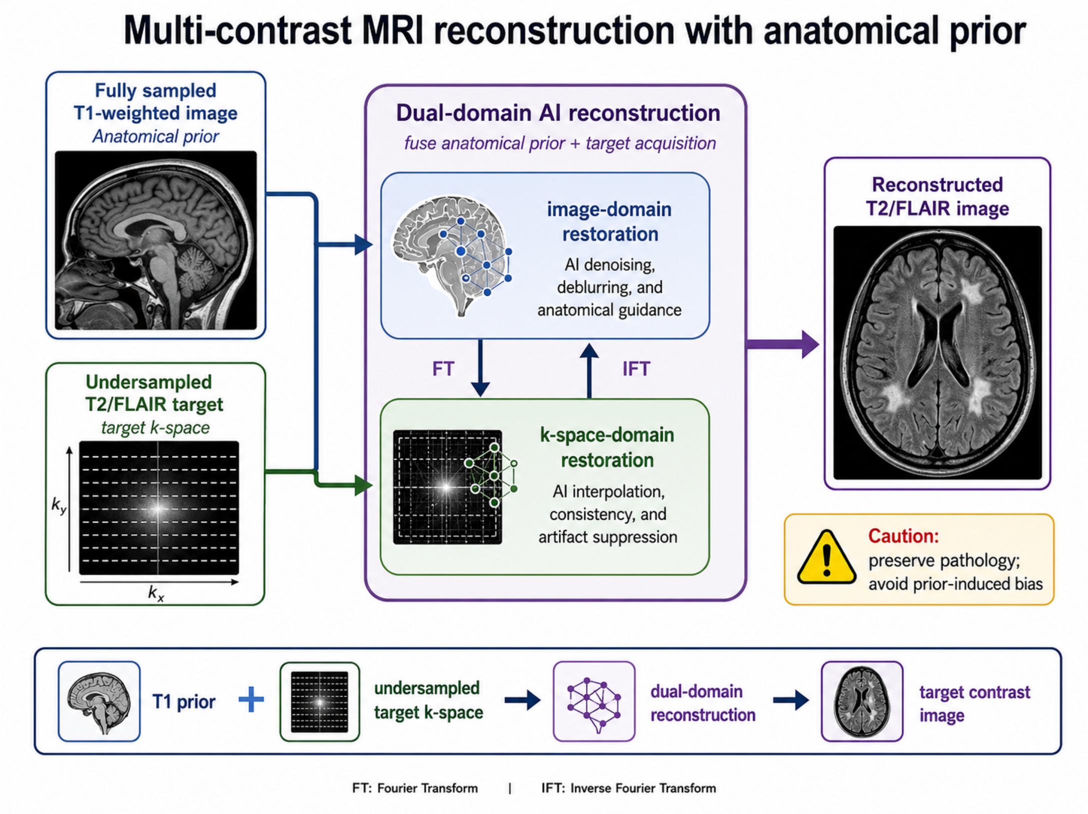

Figure 7: Multi-contrast MRI reconstruction with anatomical prior information. The figure should show a fully sampled T1-weighted image acting as an anatomical prior for reconstructing an undersampled longer protocol such as T2 or FLAIR.

4.9 AI-Guided Acquisition and Sampling

AI can also operate at the acquisition stage by deciding which k-space samples to collect. Active sampling and reinforcement-learning approaches treat this process as a sequential decision problem. L2SR extends this idea by formulating MRI sampling and reconstruction as a joint optimization problem and introducing a sparse-reward POMDP with alternating training of sampler and reconstructor to reduce mismatch between learned sampling policies and reconstruction models [7, 23].

Figure 8 shows the adaptive-acquisition loop, in which a learned policy selects additional k-space measurements in interaction with the evolving reconstruction.

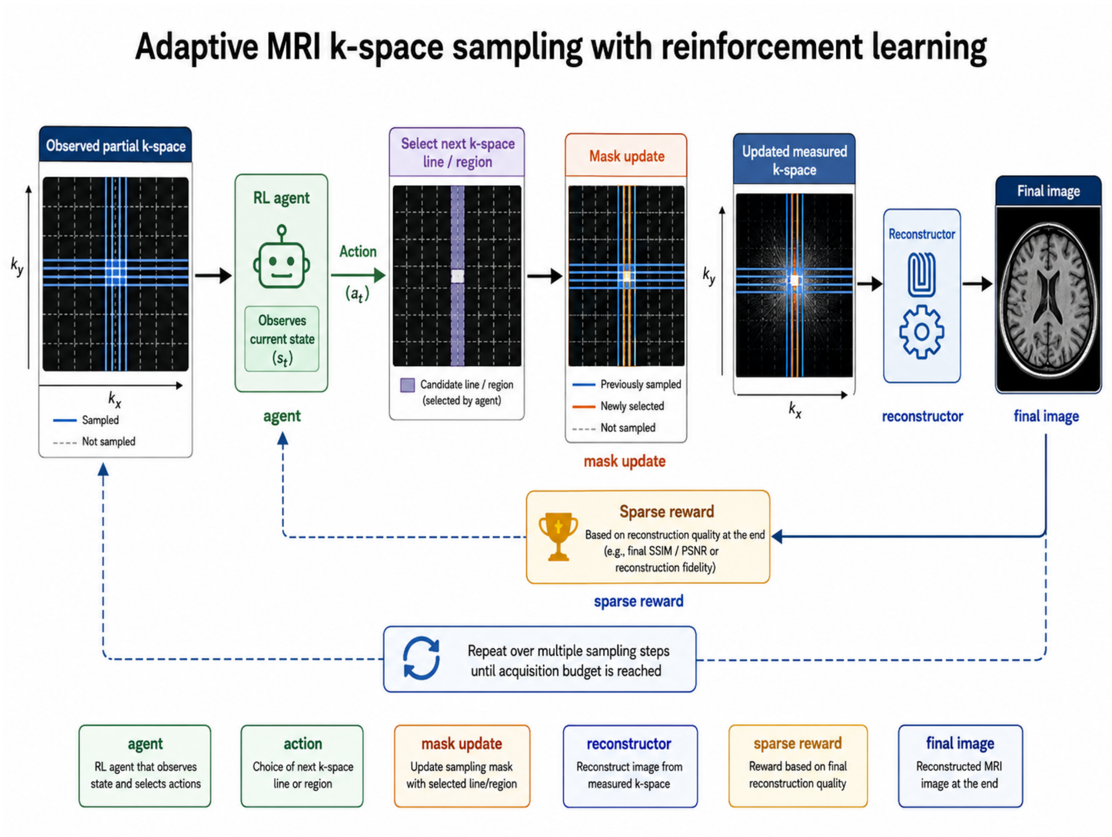

Figure 8: AI-guided acquisition and adaptive k-space sampling. The figure should show an agent selecting k-space lines sequentially, receiving reconstruction-quality feedback, and coupling the learned sampler with a learned reconstructor.

5 Benchmarks and Clinical Validation

5.1 fastMRI as a Benchmark Ecosystem

The fastMRI challenges are essential because they provide shared datasets, benchmark tasks, and clinically informed evaluation. The 2019 fastMRI challenge helped establish a common framework for machine-learning MRI reconstruction through an open competition based on raw knee MRI data and radiologist-supported evaluation [24]. The 2020 fastMRI challenge expanded the setting to brain MRI, emphasized pathology depiction, used a purely multi-coil reconstruction setting, and introduced a transfer track to test generalization across vendors [12].

5.2 Quantitative Metrics Versus Clinical Assessment

Metrics such as PSNR and SSIM are useful but insufficient on their own. The 2020 fastMRI challenge used SSIM for quantitative ranking and finalist selection, but relied on radiologist evaluation focused on the quality of pathology depiction for qualitative assessment. This supports the clinical claim that reconstruction quality must ultimately be judged by whether clinically meaningful information is preserved, not only by pixel-level similarity or global image metrics [12].

Figure 9 organizes these validation requirements into a practical evaluation pipeline that combines quantitative metrics with clinical and generalization-oriented assessment.

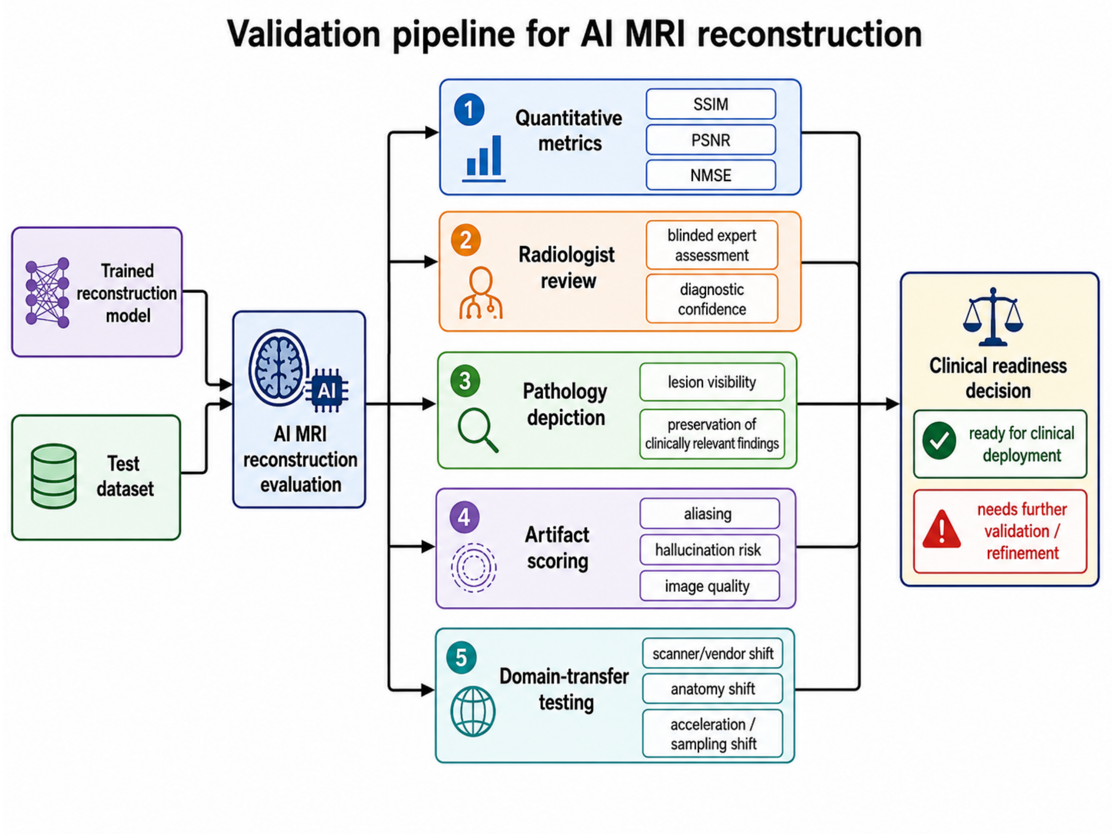

Figure 9: Validation pipeline for AI MRI reconstruction. The figure should show quantitative metrics, radiologist review, pathology depiction, artifact scoring, and vendor/domain transfer evaluation.

6 Reliability, Robustness, and Clinical Trust

6.1 Instability and Hallucination Risk

A serious report on AI MRI reconstruction must include risk. Antun et al. argue that deep-learning image reconstruction can be unstable under small perturbations, small structural changes, and changes in the sampling pattern. In medical imaging, such failures matter because they may produce misleading artifacts that are difficult to recognize as nonphysical, or suppress diagnostically relevant structures [15].

6.2 A Balanced Interpretation of Robustness

The later robustness study complicates the simple claim that deep learning is uniquely fragile. It compares trained neural networks, untrained neural networks, and classical sparsity-based methods, finding that vulnerability to small adversarial perturbations is not exclusive to trained deep learning. In the studied setting, the strongest deep-learning methods also performed well under realistic distribution shifts and in recovering small diagnostically relevant details [16]. This supports a balanced conclusion: AI reconstruction requires careful scrutiny, but classical methods likewise require explicit robustness evaluation.

6.3 Domain Shift and Generalization

Generalization is a major barrier to clinical deployment. MRI varies by anatomy, contrast, scanner vendor, field strength, coil configuration, sampling pattern, acceleration factor, and other acquisition- and patient-specific factors [1]. The systematic evaluation paper shows that at high acceleration, especially $R = 8$, alignment between training and test data can matter more than architectural choice [8]. The 2020 fastMRI transfer track reinforces this point by explicitly evaluating models on vendors outside the training set [12].

6.4 Clinical Trust Requirements

Clinical translation requires more than visually appealing reconstructions. It requires preservation of diagnostically relevant pathology, robustness under distribution shift, characterization of reconstruction artifacts and instability modes, radiologist reader studies, prospective clinical validation, mechanisms for identifying reconstruction failures, and integration into practical clinical workflows. These requirements are motivated jointly by the fastMRI challenge design, the instability and robustness literature, systematic evaluations of generalization, and broader reviews of clinical translation in MRI reconstruction [1, 8, 12, 15, 16].

Figure 10 summarizes representative failure modes that should be considered when evaluating clinical trustworthiness, including perturbation sensitivity, missed detail, hallucinated structure, and domain mismatch.

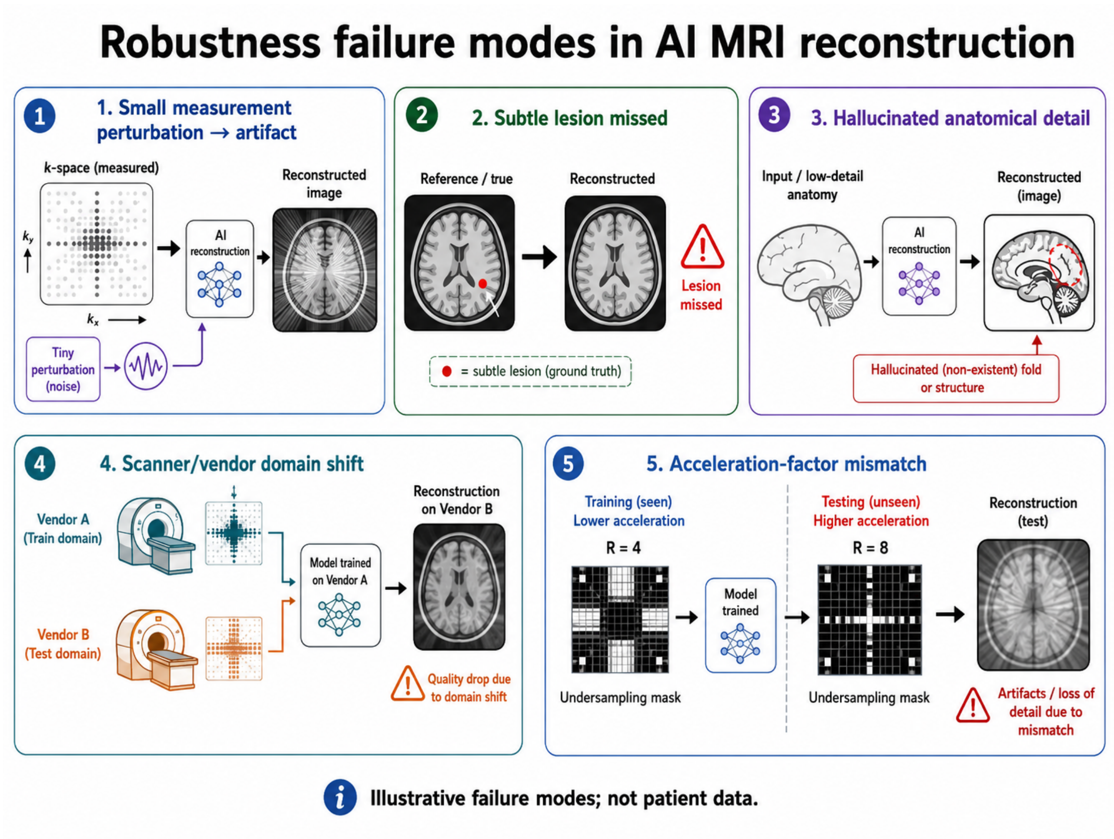

Figure 10: Robustness and failure modes. The figure should illustrate perturbation sensitivity, missed small structures, hallucinated anatomy, scanner/vendor domain shift, and acceleration-factor mismatch as distinct validation challenges.

7 Brief Comparison with CT and PET

CT and PET are useful comparison points, but they should not dominate this MRI-centered report. CT reconstruction is commonly discussed in relation to radiation dose, sparse-view acquisition, and noise texture. PET reconstruction and denoising often focus on low-count statistics, tracer uptake, and anatomical priors from CT or MR.

The PET studies considered here emphasize a different problem structure. One unsupervised PET denoising approach uses same-patient CT or MR prior images as network input and the noisy PET image itself as the training target, while another PET denoising approach uses simulation-based pretraining, real-data fine-tuning, and perceptual loss to address limited labeled data and improve detail preservation [25, 26]. MRI reconstruction differs because its central structure is governed by k-space undersampling, Fourier encoding, multi-coil acquisition, protocol-specific contrast information, and explicit data consistency with acquired measurements [1, 4].

Figure 11 places this distinction in a broader modality comparison by contrasting MRI reconstruction with characteristic CT and PET AI reconstruction tasks.

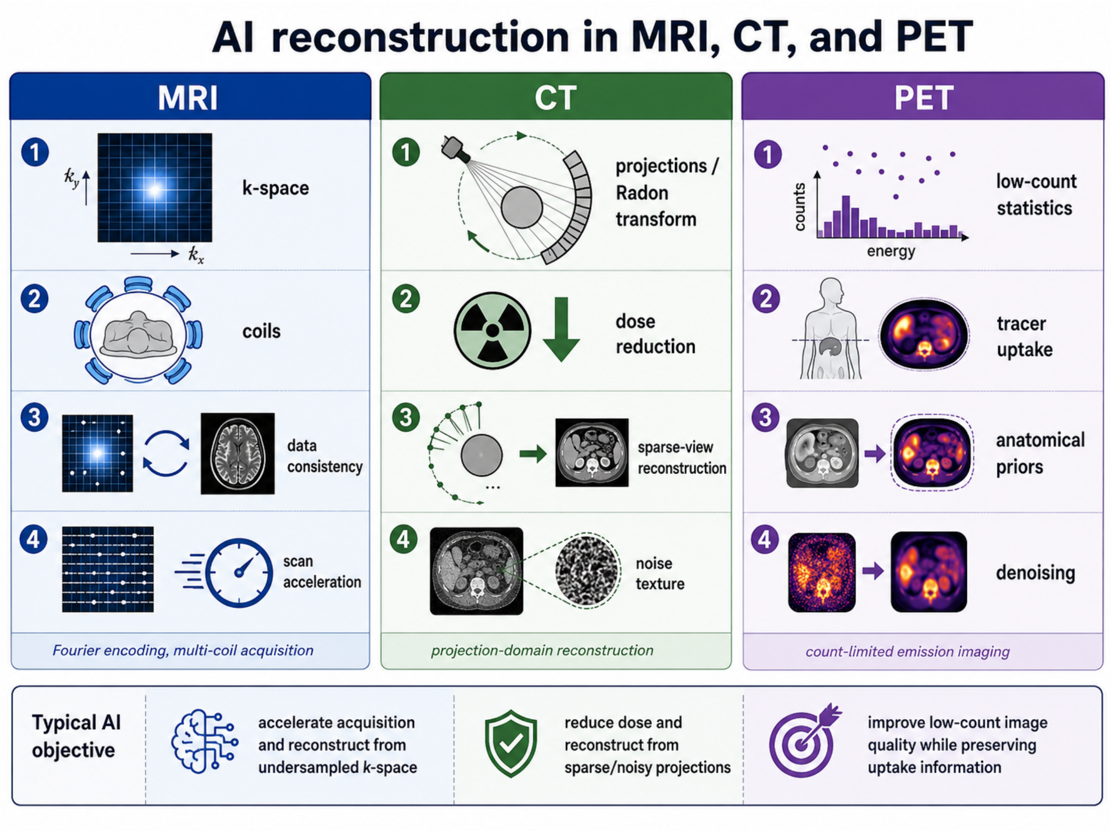

Figure 11: Comparison of AI reconstruction tasks across MRI, CT, and PET. The figure should contrast MRI k-space/coil/data-consistency reconstruction with CT dose/sparse-view reconstruction and PET low-count/anatomical-prior denoising.

8 Future Directions

8.1 Physics-Aware Hybrid Systems

A strongly supported future direction is physics-aware AI for MRI reconstruction. This includes unrolled networks, recurrent variational methods, dual-domain architectures, optimized compressed-sensing hybrids, and methods that explicitly preserve consistency with measured k-space data [1, 6, 8–10].

8.2 Better Self-Supervision

Self-supervised and unsupervised MRI reconstruction will remain important because fully sampled reference data are costly and, in some settings, difficult or impractical to obtain under realistic clinical constraints. MRI-specific approaches therefore seek to learn from undersampled acquisitions while still incorporating the acquisition model and preserving consistency with measured data [20, 21]. More general denoising work on learning without clean paired targets, such as Recorrupted-to-Recorrupted and unpaired deep image denoising, can contribute useful conceptual ideas, but it should remain background rather than the methodological center of an MRI reconstruction discussion [27, 28].

8.3 Adaptive Acquisition

Adaptive acquisition is an especially promising direction because it addresses what should be measured, rather than merely repairing undersampled data after the fact. Reinforcement-learning approaches formulate k-space acquisition as a sequential decision problem, while joint sampler–reconstructor methods extend this idea by optimizing acquisition and reconstruction together [7, 23].

8.4 Evaluation Beyond SSIM

Future validation should combine quantitative metrics, radiologist assessment, pathology preservation, domain-shift testing, robustness checks, and prospective clinical studies [12, 15, 16].

9 Conclusion

AI has changed MRI reconstruction by introducing trainable priors, learned optimization, dual-domain reconstruction, and adaptive acquisition into a problem historically shaped by parallel imaging and compressed sensing. The most credible direction is not unconstrained image generation but data-consistent, physics-informed, multi-coil, clinically evaluated reconstruction. Classical methods remain important both as baselines and as conceptual foundations. The literature reviewed here suggests that future MRI reconstruction will be hybrid: combining the reliability of physical models with the flexibility of learned methods [6–10, 12, 13, 23].

Bibliography

[1] Javier Montalt-Tordera, Vivek Muthurangu, Andreas Hauptmann, and Jennifer Anne Steeden. Machine learning in magnetic resonance imaging: Image reconstruction. *Physica Medica*, 83:79–87, 2021. doi: 10.1016/j.ejmp.2021.02.020.

[2] Arghya Pal and Yogesh Rathi. A review and experimental evaluation of deep learning methods for MRI reconstruction. *Machine Learning for Biomedical Imaging*, 1(March 2022 issue):1–50, 2022. doi: 10.59275/j.melba.2022-3g12.

[3] Chang Min Hyun, Hyeongseok P. Kim, Sung Min Lee, Sung Lee, and Jin Keun Seo. Deep learning for undersampled MRI reconstruction. *Physics in Medicine & Biology*, 63(13):135007, 2018. doi: 10.1088/1361-6560/aac71a.

[4] Florian Knoll, Kerstin Hammernik, Chi Zhang, Steen Moeller, Thomas Pock, Daniel K. Sodickson, and Mehmet Akçakaya. Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues. *IEEE Signal Processing Magazine*, 37 (1):128–140, 2020. doi: 10.1109/MSP.2019.2950640.

[5] Christopher M. Sandino, Joseph Y. Cheng, Fei Chen, Morteza Mardani, John M. Pauly, and Shreyas S. Vasanawala. Compressed sensing: From research to clinical practice with deep neural networks. *IEEE Signal Processing Magazine*, 37(1):111–127, 2020. doi: 10.1109/MSP.2019.2950433.

[6] Hongyi Gu, Burhaneddin Yaman, Steen Moeller, Jutta Ellermann, Kamil Uğurbil, and Mehmet Akçakaya. Revisiting ℓ_1-wavelet compressed-sensing MRI in the era of deep learning. *Proceedings of the National Academy of Sciences*, 119(33):e2201062119, 2022. doi: 10.1073/pnas.2201062119.

[7] Luis Pineda, Soumya Basu, Adriana Romero, Roberto Calandra, and Michal Drozdzal. Active MR k-space sampling with reinforcement learning. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pages 23–33. Springer, 2020. doi: 10.1007/978-3-030-59713-9_3.

[8] Kerstin Hammernik, Jo Schlemper, Chen Qin, Jinming Duan, Ronald M. Summers, and Daniel Rueckert. Systematic evaluation of iterative deep neural networks for fast parallel MRI reconstruction with sensitivity-weighted coil combination. *Magnetic Resonance in Medicine*, 86(4):1859–1872, 2021. doi: 10.1002/mrm.28827.

[9] George Yiasemis, Jan-Jakob Sonke, Clarisa Sánchez, and Jonas Teuwen. Recurrent variational network: A deep learning inverse problem solver applied to the task of accelerated MRI reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 732–741, 2022.

[10] Bo Zhou and S. Kevin Zhou. DuDoRNet: Learning a dual-domain recurrent network for fast MRI reconstruction with deep T1 prior. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4273–4282. IEEE, 2020. doi: 10.1109/CVPR42600.2020.00433.

[11] Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, Jutta Ellermann, Kamil Uğurbil, and Mehmet Akçakaya. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magnetic Resonance in Medicine*, 84(6):3172–3191, 2020. doi: 10.1002/mrm.28378.

[12] Matthew J. Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, Simon Arberet, Dominik Nickel, Zaccharie Ramzi, Philippe Ciuciu, Jean-Luc Starck, Jonas Teuwen, Dimitrios Karkalousos, Chi Zhang, Anuroop Sriram, Zhengnan Huang, Nafissa Yakubova, Yvonne W. Lui, and Florian Knoll. Results of the 2020 fastMRI challenge for machine learning MR image reconstruction. *IEEE Transactions on Medical Imaging*, 40(9):2306–2317, 2021. doi: 10.1109/TMI.2021.3075856.

[13] Anuroop Sriram, Jure Zbontar, Tullie Murrell, C. Lawrence Zitnick, Aaron Defazio, and Daniel K. Sodickson. GrappaNet: Combining parallel imaging with deep learning for multi-coil MRI reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14315–14322, 2020.

[14] Mahyar Safari, Zaid Eidex, Chia-Wei Chang, R. L. J. Qiu, and Xiaofeng Yang. Fast MRI reconstruction using deep learning-based compressed sensing: A systematic review, 2024.

[15] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C. Hansen. On instabilities of deep learning in image reconstruction and the potential costs of AI. *Proceedings of the National Academy of Sciences*, 117(48):30088–30095, 2020. doi: 10.1073/pnas.1907377117.

[16] Mohammad Zalbagi Darestani, Akshay S. Chaudhari, and Reinhard Heckel. Measuring robustness in deep learning based compressive sensing. In *Proceedings of the 38th International Conference on*

Machine Learning, volume 139 of *Proceedings of Machine Learning Research*, pages 2433–2444, 2021.

[17] Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, and Dong Liang. Accelerating magnetic resonance imaging via deep learning. In *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, pages 514–517. IEEE, 2016. doi: 10.1109/ISBI.2016.7493320.

[18] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony N. Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for MR image reconstruction. *IEEE Transactions on Medical Imaging*, 37(2):491–503, 2018. doi: 10.1109/TMI.2017.2760978.

[19] Jian Lv, Chen Wang, and Guang Yang. PIC-GAN: A parallel imaging coupled generative adversarial network for accelerated multi-channel MRI reconstruction. *Diagnostics*, 11(1):61, 2021. doi: 10.3390/diagnostics11010061.

[20] Guohua Zeng, Yao Guo, Jiaming Zhan, Zhi Wang, Zhen Lai, Xiaoyan Du, Xiaobo Qu, and Dong Guo. A review on deep learning MRI reconstruction without fully sampled k-space. *BMC Medical Imaging*, 21:195, 2021. doi: 10.1186/s12880-021-00727-9.

[21] Kuang Gong, Paul K. Han, Georges El Fakhri, Chuan Ma, and Quanzheng Li. Arterial spin labeling MR image denoising and reconstruction using unsupervised deep learning. *NMR in Biomedicine*, 35 (4):e4224, 2022. doi: 10.1002/nbm.4224.

[22] Lei Xiang, Yutong Chen, Wei Chang, Yiqiang Zhan, Weili Lin, Qian Wang, and Dinggang Shen. Deep-learning-based multi-modal fusion for fast MR reconstruction. *IEEE Transactions on Biomedical Engineering*, 66(7):2105–2114, 2019. doi: 10.1109/TBME.2018.2883958.

[23] Peiyu Yang and Bin Dong. L2SR: Learning to sample and reconstruct for accelerated MRI via reinforcement learning. *Inverse Problems*, 40(5):055015, 2024. doi: 10.1088/1361-6420/ad3b34.

[24] Florian Knoll, Tullie Murrell, Anuroop Sriram, Nafissa Yakubova, Jure Zbontar, Michael Rabbat, Aaron Defazio, Matthew J. Muckley, Daniel K. Sodickson, C. Lawrence Zitnick, and Michael P. Recht. Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge. *Magnetic Resonance in Medicine*, 84(6):3054–3070, 2020. doi: 10.1002/mrm.28338.

[25] Jian Cui, Kuang Gong, Ning Guo, Chaohui Wu, Xiaotong Meng, Kang Kim, Kai Zheng, Zhi Wu, Long Fu, Bin Xu, Zhi Zhu, Jie Tian, Hong Liu, and Quanzheng Li. PET image denoising using unsupervised deep learning. *European Journal of Nuclear Medicine and Molecular Imaging*, 46 (13):2780–2789, 2019. doi: 10.1007/s00259-019-04468-4.

[26] Kuang Gong, Jian Guan, Chao-Chieh Liu, and Jinyi Qi. PET image denoising using a deep neural network through fine tuning. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 3(2): 153–161, 2019. doi: 10.1109/TRPMS.2018.2877644.

[27] Tongyao Pang, Huan Zheng, Yuhui Quan, and Hui Ji. Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2043–2052, 2021.

[28] Xiaohe Wu, Ming Liu, Yue Cao, Dongwei Ren, and Wangmeng Zuo. Unpaired learning of deep image denoising. In *Computer Vision – ECCV 2020*, volume 12349 of *Lecture Notes in Computer Science*, pages 352–368. Springer, 2020. doi: 10.1007/978-3-030-58548-8_21.