

Mind the Detail: Uncovering Clinically Relevant Image Details in Accelerated MRI with Semantically Diverse Reconstructions

Jan Nikolas Morshuis¹, Christian Schlarmann¹, Thomas Küstner^{1,2},
Christian F. Baumgartner^{1,3*}, and Matthias Hein^{1*}

¹ University of Tübingen, Germany

² University Hospital of Tübingen, Germany

³ University of Lucerne, Switzerland

`nikolas.morshuis@uni-tuebingen.de`

*Shared last authorship

Abstract. In recent years, accelerated MRI reconstruction based on deep learning has led to significant improvements in image quality with impressive results for high acceleration factors. However, from a clinical perspective image quality is only secondary; much more important is that all clinically relevant information is preserved in the reconstruction from heavily undersampled data. In this paper, we show that existing techniques, even when considering resampling for diffusion-based reconstruction, can fail to reconstruct small and rare pathologies, thus leading to potentially wrong diagnosis decisions (false negatives). To uncover the potentially missing clinical information we propose “Semantically Diverse Reconstructions” (SDR), a method which, given an original reconstruction, generates novel reconstructions with enhanced semantic variability while all of them are fully consistent with the measured data. To evaluate SDR automatically we train an object detector on the fastMRI+ dataset. We show that SDR significantly reduces the chance of false-negative diagnoses (higher recall) and improves mean average precision compared to the original reconstructions. The code is available on <https://github.com/NikolasMorshuis/SDR>.

Keywords: MRI Reconstruction · Detection · Uncertainty

1 Introduction

Accelerated MRI reconstruction from undersampled k-space data is an inverse problem with infinitely many solutions. However, it is common to reconstruct only a single solution, thereby ignoring potentially different yet still plausible reconstructions. Deep-learning-based approaches, including novel techniques based on diffusion models, have led to impressive reconstruction results in terms of image quality even for high acceleration factors. On the other hand, in clinical practice, some pathological tissue changes can be small and may occur rarely, making them potentially underrepresented in the training data. As a result,




Fig. 1: **Method overview.** At high acceleration factors (12x shown here), naive MRI reconstruction suffers from artifacts (left). Recent reconstruction methods like DDIP [4] can recover image details and generate multiple reconstructions, but these often lack diversity, potentially leading to missed pathologies (top right). Our proposed SDR approach produces semantically diverse reconstructions, helping detect pathologies that might otherwise be missed (bottom right).

overconfident reconstruction models may favor more likely solutions of healthy tissue if consistent with the measurement data, potentially leading to missed pathologies. Missing clinically relevant pathologies in the reconstruction can be a major obstacle to the application of accelerated MRI in clinical practice.

While some recent reconstruction methods such as diffusion-models [3,4,11] are capable of generating multiple different solutions by repeated sampling, it has been shown that these models tend to sample the most likely solutions frequently but only rarely sample less likely solutions that might still be diagnostically relevant [5] (see Fig. 1 top right). Moreover, sampling multiple reconstructions can be time-consuming as it requires repeating the denoising process.

To address this problem, we propose a technique to generate semantically diverse reconstructions (SDR) that are fully consistent with the measured k-space data. Our SDR approach can be applied post-hoc to initial reconstructions obtained using any reconstruction method, producing a set of alternative solutions that preserve data consistency while potentially revealing pathologies not present in the original reconstruction. This offers the unique possibility to study a diverse set of data-consistent solutions and their inherent uncertainty. Understanding the inherent uncertainty can improve the safety of the reconstruction process because the risk of a missed pathology and therefore a false negative diagnosis is reduced. If high uncertainty prevents a confident diagnosis from being made, additional measurements can be acquired. Our method generates these alternative solutions in ≈ 3 seconds per image, which is faster than the sampling time of diffusion models that can require several minutes per image [11]. This

fast inference time could potentially enable an application of our method during the scanning process in the future.

Our technique relies on identifying candidate regions in reconstructed images and modifying them to alter their semantic content while ensuring consistency with the measured k-space data. To generate semantically diverse reconstructions, we develop an adversarially robust version of the detection algorithm ViTDet [14], leveraging techniques from [20] for the ViT-Backbone. We then use the features of candidate regions to maximize the semantic distance from the existing reconstructions (see Fig. 1 bottom right). This is motivated by the observation that the Euclidean distance in the feature space of adversarially robust vision encoders is a good semantic metric and that maximizing distance in this feature space leads to semantically meaningful changes [6]. We explore two versions of our method: i) an automatic selection of regions to be optimized using the proposal network of the object detector, ii) manual annotation of a candidate image region by the human user, e.g. a doctor.

In experiments on the fastMRI+ dataset [27], we demonstrate that our method is able to produce a semantically diverse set of possible MRI reconstructions that capture more possible pathologies compared to alternative sampling techniques based on diffusion models [4] and the E2E-Varnet [22]. We quantitatively verify this by showing that pathology recall – obtained with a separate Faster-RCNN detector [18] – outperforms the baseline techniques, while mean average precision is not negatively affected.

2 Related Work

The generation of multiple solutions in deep-learning-based MRI reconstructions is a relatively recent area of research. Morshuis et al. [17] proposed a method to produce diverse MRI-reconstructions corresponding to minimum and maximum segmentations in a downstream task. This approach estimates the segmentation uncertainty due to the ill-posed nature of the reconstruction problem. Küstner et al. [13] trained models with different initial seeds to generate multiple reconstructions, which were then used to estimate reconstruction uncertainty. They observed that uncertainty was highest in image regions containing pathologies. Recent studies have demonstrated the strong performance of diffusion models in MR reconstruction [11,3,4,25]. Apart from producing high-quality reconstructions, these techniques can also be used to generate multiple solutions by repeated sampling from the posterior distribution of images that are consistent with the measurement data.

Only few works have focused on pathology detection in MRI [10,26]. Other work has used the pathology annotations provided by some datasets [27,7] to compare reconstruction metrics in these relevant areas [13,7] or to show that measurement noise can lead to changes in reconstructions in these areas [16]. To our knowledge, ours is the first work to consider pathology detection in accelerated MRI reconstruction utilizing multiple reconstructions.

3 Method

3.1 Background and Notation

The goal of MRI reconstruction is to find a reconstructed image $\hat{\mathbf{x}}$ that is consistent with the measured k-space data \mathbf{y} subject to the following equation:

$$\mathbf{y} = \mathbf{A}\hat{\mathbf{x}} \quad \text{with} \quad \mathbf{y} \in \mathbb{C}^m, \quad \mathbf{A} \in \mathbb{C}^{m \times n}, \quad \hat{\mathbf{x}} \in \mathbb{C}^n, \quad (1)$$

where $\mathbf{A} = MFS$ with the coil sensitivity maps S , the Fourier transform \mathcal{F} and the undersampling mask M . MRI acquisitions are often accelerated by only measuring a fraction of the k-space data \mathbf{y} .

The fully sampled image, which would have been obtained by measuring the entire k-space, contains a (possibly empty) set of pathologies denoted as $\mathcal{S} = \{s_1, s_2, \dots\}$. However, due to image degradation caused by acceleration, a different set of pathologies, $\mathcal{S}' = \{s'_1, s'_2, \dots\}$, is visible in the reconstructed image $\hat{\mathbf{x}}$. Ideally, \mathcal{S} and \mathcal{S}' should be identical to ensure accurate diagnosis. However, there are infinitely many reconstructions that satisfy Eq. 1 with no guarantee that they correctly reconstruct pathologies. As a result, \mathcal{S}' may fail to include all pathologies present in the patient, particularly for high acceleration factors.

In this work, we aim to generate a set of reconstructions, $\mathcal{X} = \{\hat{\mathbf{x}}^{(1)}, \dots, \hat{\mathbf{x}}^{(N_{rec})}\}$, such that when considered jointly, the reconstructed images capture a set of pathologies \mathcal{S}' that includes as many of the ground truth pathologies \mathcal{S} as possible (few false negatives) while not hallucinating pathologies (few false positives).

3.2 Generating Diverse Solutions Maximizing Semantic Distance

Our key contribution is a method which, given an initial reconstruction $\hat{\mathbf{x}}^{(1)}$, can efficiently generate a set of reconstructions \mathcal{X} that are semantically diverse in the regions where pathologies are likely to occur. To achieve this, we iteratively generate additional reconstructions $\hat{\mathbf{x}}^{(i)}$, $i = 2, \dots, N_{rec}$, by maximizing semantic distance d to all previous reconstructions $\hat{\mathbf{x}}^{(j)}$ in an iterative procedure

$$\hat{\mathbf{x}}^{(i)} \leftarrow \mathcal{DC} \left(P_{B(\hat{\mathbf{x}}^{(1)}, r)} \left(\hat{\mathbf{x}}^{(i)} + \eta \nabla_{\hat{\mathbf{x}}^{(i)}} \sum_{j \neq i} d(\hat{\mathbf{x}}^{(i)}, \hat{\mathbf{x}}^{(j)}) \right) \right), \quad (2)$$

where η is the step size, $\nabla_{\hat{\mathbf{x}}^{(i)}}$ is the gradient operator, $P_{B(\hat{\mathbf{x}}^{(1)}, r)}$ is the projection onto the ℓ_2 -ball of radius r around $\hat{\mathbf{x}}^{(1)}$, and \mathcal{DC} is the projection operator mapping the reconstruction $\hat{\mathbf{x}}^{(i)}$ on the data-consistent space. We ensure data-consistency by replacing the k-space data with the original measured data at the locations where a measurement was acquired. To prevent potential artifacts from occurring, we run M-Step Conjugate Gradient before, minimizing the distance of the predicted k-space to the measurement. Furthermore, the projection $P_{B(\hat{\mathbf{x}}^{(1)}, r)}$ ensures that $\|\hat{\mathbf{x}}^{(i)} - \hat{\mathbf{x}}^{(1)}\|_2 \leq r$, where r is the allowed perturbation radius set to $r = 3$ in our experiments. The value has been chosen, as it allows for meaningful perturbations while keeping the number of new false positives low.

Input: $\mathbf{y}, \hat{\mathbf{x}}^{(1)}, f_\phi, RPN, N_{opt}, N_{rec}, r$
Output: $\mathcal{X} = \{\hat{\mathbf{x}}^{(1)}, \dots, \hat{\mathbf{x}}^{(N_{rec})}\}$
 $\mathcal{B} = \{b_1, b_2, \dots\} \leftarrow RPN(\hat{\mathbf{x}}^{(1)})$ or manually ; // Generate Proposal Boxes
for $i \leftarrow 1$ **to** N_{rec} **do**
 $\boldsymbol{\varepsilon}^{(i)} \sim \mathcal{N}(0, \sigma^2)$ with $\boldsymbol{\varepsilon}^{(1)} = 0$; // Keep init. recon.
 $\hat{\mathbf{x}}^{(i)} \leftarrow P_{B(\hat{\mathbf{x}}^{(1)}, r)}(\hat{\mathbf{x}}^{(1)} + \boldsymbol{\varepsilon}^{(i)})$, where $B(\hat{\mathbf{x}}^{(1)}, r) = \{\mathbf{x} \mid \|\mathbf{x} - \hat{\mathbf{x}}^{(1)}\|_2 \leq r\}$,
 $\hat{\mathbf{x}}^{(i)} \leftarrow \mathcal{DC}(\hat{\mathbf{x}}^{(i)})$; // Data consistency
 $\mathbf{a}^{(i)} \leftarrow f_\phi(\hat{\mathbf{x}}^{(i)}, \mathcal{B})$; // Get box-features
end
for $k \leftarrow 0$ **to** N_{opt} **do**
 for $i \leftarrow 2$ **to** N_{rec} **do**
 $d^{(i)} = \sum_{j \neq i} \|\mathbf{a}^{(i)} - \mathbf{a}^{(j)}\|_2$; // Calculate feature distances
 $\hat{\mathbf{x}}^{(i)} \leftarrow P_{B(\hat{\mathbf{x}}^{(1)}, r)}(\hat{\mathbf{x}}^{(i)} + \eta \nabla_{\hat{\mathbf{x}}^{(i)}} d^{(i)})$
 $\hat{\mathbf{x}}^{(i)} \leftarrow \mathcal{DC}(\hat{\mathbf{x}}^{(i)})$; // Data consistency
 $\mathbf{a}^{(i)} \leftarrow f_\phi(\hat{\mathbf{x}}^{(i)}, \mathcal{B})$; // Get box-features
 end
end

Algorithm 1: Generating semantically diverse reconstructions by optimizing the distance between box features $\mathbf{a}^{(i)}$ obtained by the robust ViTDet.

To measure semantic distance between images, we utilize the features of a ViTDet detection network f_ϕ [14], which we have finetuned on the fastMRI+ data. Relevant proposal bounding boxes $\mathcal{B} = \{b_1, b_2, \dots\}$ can be selected using methods described in Sec. 3.4. Given the box-features $\mathbf{a}^{(i)} = f_\phi(\hat{\mathbf{x}}^{(i)}, \mathcal{B})$ we define

$$d(\hat{\mathbf{x}}^{(i)}, \hat{\mathbf{x}}^{(j)}) = \sum_{b \in \mathcal{B}} \|f_\phi(\hat{\mathbf{x}}^{(i)}, b) - f_\phi(\hat{\mathbf{x}}^{(j)}, b)\|_2. \quad (3)$$

The proposed algorithm is summarized in Alg. 1.

3.3 Adversarially Robust Backbone

Our application relies on gaining semantically meaningful gradients from the box-feature encoder f_ϕ in Eq. (3). While gradients of deep learning models are generally not interpretable [15], it is well-known that adversarially robust models possess semantically more meaningful gradients [24, 19, 1, 9, 6]. To obtain an adversarially robust detection model, we leverage the recently proposed adversarial fine-tuning method of [20]. We apply this method to the ViT-backbone g_ϕ of the ViTDet. The ViT-backbone also serves as the backbone of the box-feature encoder $f_\phi = h_\phi \circ g_\phi$. The method ensures that the backbone features, and thus also the box features, remain stable under small adversarial perturbations. The training objective is

$$\min_{\phi} \sum_i \max_{\|\boldsymbol{\delta}\|_2 \leq r} \|g_\phi(\mathbf{x}_i + \boldsymbol{\delta}) - g_{\phi_{\text{org}}}(\mathbf{x}_i)\|_2^2, \quad (4)$$




Fig. 2: **Automatic Proposal Generation.** Given an initial reconstruction, we filter the proposal boxes given by the Region Proposal Network (RPN) of the ViTDet to only keep the 75% of boxes with the highest score. We then apply Non-Maximum Suppression (NMS) on the remaining proposals. The remaining proposal boxes are then used for optimization with SDR-A.

where $g_{\phi_{\text{org}}}$ is the frozen original backbone, r is the adversarial perturbation budget, and \mathbf{x}_i are unlabeled training images. In contrast to [20], we use the ℓ_2 -norm instead of the ℓ_∞ -norm to bound the perturbation budget, as this generally yields better generative capabilities [1]. We set r to 10 and train for 20,000 steps on ImageNet data at resolution 640 using a batch size of 16. While it would be natural to train on fastMRI images, we have observed a stronger performance when training on a much larger set of diverse natural images.

3.4 Box-Proposal Selection

We introduce two approaches for identifying the box regions to optimize, which are motivated by two different use-cases.

With **Manual (SDR-M) Proposal Box Generation** we address the scenario where a user suspects a pathology in a specific area and manually draws a bounding box. We simulate user annotations by perturbing the ground-truth boxes in size (75%-125%) and position (up to 25% of width or height).

With **Automatic (SDR-A) Proposal Box Generation** we address the case when no patient-specific information on the location of pathologies is available. In this scenario, we use the Region Proposal Network (RPN) of the ViTDet. We filter out the 25% of boxes that have the lowest probability, as these are often in the background of the MRI image. We then utilize Non-Maximum-Suppression (NMS) to remove boxes with significant overlap ($\text{IoU} > 0.05$), see Figure 2.

After applying either method, we get a set of proposal boxes $\mathcal{B} = \{b_1, b_2, \dots\}$ in which we optimize the reconstructions to be semantically different.

4 Data, Evaluation and Results

Assessing clinical relevance of reconstructions would ideally require input from clinicians. As it is difficult to do such an evaluation at scale, instead, we evaluate by training an object detector and use it as a proxy for human annotators.

Data and Training: Open datasets of MRI reconstruction data are rare [12,7,21,23]




Fig. 3: **Combining predictions of multiple reconstructions.** We combine overlapping detection boxes and average their probability score across reconstructed images to compute the mAP of our detections.

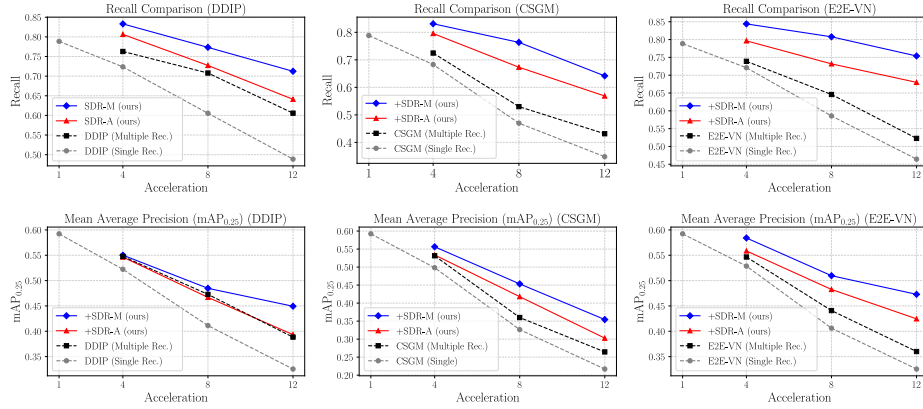


Fig. 4: **Evaluation of reconstruction methods.** Both SDR-A and SDR-M outperform single and multiple reconstructions using the original reconstruction techniques (DDIP, CSGM, E2E-VN) in terms of recall (better detection of pathologies) while achieving similar or better $mAP_{0.25}$. Thus SDR generates clinically relevant reconstructions outperforming resampling.

and only a few provide annotated pathologies [27,7]. For our experiments we use the fastMRI+ [27] knee data, as it provides enough annotated slices to train large detection networks. We split the fastMRI+ validation set subject-wise into test and validation set. For evaluation, we considered the 10 most frequent pathology classes in fastMRI+, as the remaining classes have not enough test samples.

Object Detector for Evaluation: To ensure an objective evaluation independent of the ViTDet used for generating the reconstructions of SDR, we trained a separate Faster-RCNN detection network [18] which we use for evaluation.

Combining detections of multiple images: Given a set of potential pathology detections $\mathcal{S}' = \{s'_1, s'_2, \dots\}$ that was extracted from a set of reconstructions \mathcal{X} , we need to remove near duplicates and estimate the probability of each bounding box for the calculation of mean average precision (mAP). To remove near duplicates, we used NMS to eliminate overlapping bounding boxes (Intersection-over-Union (IoU) ≥ 0.25) of the same class. To estimate the probability scores

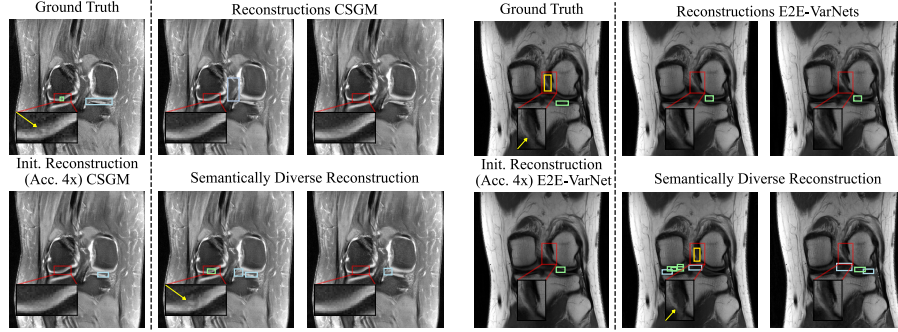


Fig. 5: **Qualitative Results** The highlighted box in the ground truth shows a cartilage thickness loss (left) and a high grade ligament sprain (right). The initial reconstructions are often too smooth, preventing the detection of the pathologies. The semantic changes of SDR-M (left) and SDR-A (right) reveal the pathological changes. Yellow arrows indicate pathologies.

of the remaining boxes, we averaged the probability for a given bounding box s'_i across all reconstructed images and set the probability score to 0 for reconstructions that did not contain the bounding box.

Reconstruction Techniques: We analyzed the benefits of SDR with three base reconstruction methods: the diffusion-based DDIP [4], and CSGM [11] approaches, as well as the E2E-VarNet [22]. We produce 3 total reconstructions for each image using the baselines as well as SDR, i.e. SDR ($N_{rec} = 3$). While the diffusion models can generate multiple reconstructions by design, we trained three separate E2E-VarNets with three different random seeds for each acceleration to reconstruct three different images during inference.

Performance Measures: We measured the methods’ ability to produce sets of reconstructions that contain the ground truth pathologies using the recall, where a recall of 1 indicates no missed pathologies. As more reconstructions might lead to more false positives, we also measured mean Average Precision ($mAP_{0.25}$) [8]. For both metrics, we consider an IoU of 0.25 as correct detection. Small IoU thresholds are commonly used in medical detection tasks [2] to account for uncertain border definitions and the typically small size of pathologies.

5 Results

The quantitative results in Fig. 4 show that SDR-A and SDR-M have higher recall than the original single and multiple reconstructions. Thus SDR allows to uncover pathologies which can be missed with the original reconstruction methods. This is true for the diffusion-based reconstruction methods CSGM and DDIP as well as the often applied technique E2E-VarNet. At the same time SDR even improves $mAP_{0.25}$ and thus SDR is not simply achieving higher recall by introducing more false positive detections. We observed that it is necessary

to use a ViTDet with robust backbone for the generation in SDR, as the original ViTDet did not improve in recall or mAP compared to repeated sampling.

6 Discussion and Conclusion

We have shown that existing accelerated MRI reconstruction methods can miss pathologies and have suggested SDR to explore semantically different but fully data consistent reconstructions. The evaluation of SDR reconstructions with an object detector shows fewer missed pathologies (higher recall) while improving mAP. We believe that SDR can lead to a safer and more reliable development of image reconstructions for accelerated MRI.

The generated semantically diverse reconstructions can also be utilized by a clinician to better understand the reconstruction uncertainty, as the diverse reconstructions visualize disparities between the solutions. This would allow the clinician to acquire further measurements if necessary.

A key next step is to integrate our approach into clinical workflows and to observe how radiologists interact with semantically diverse reconstructions in practice. Furthermore, deploying our method during the scanning process could allow the system to automatically determine whether additional data acquisition is needed to further reduce semantic uncertainty.

Acknowledgments. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC number 2064/1 - Project number 390727645. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting JNM and CS.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Augustin, M., Meinke, A., Hein, M.: Adversarial robustness on in- and out-distribution improves explainability. In: ECCV (2020)
2. Baumgartner, M., Jäger, P.F., Isensee, F., Maier-Hein, K.H.: nndetection: A self-configuring method for medical object detection. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) MICCAI 2021. pp. 530–539. Springer International Publishing, Cham (2021)
3. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated mri. *Medical Image Analysis* p. 102479 (2022)
4. Chung, H., Ye, J.C.: Deep diffusion image prior for efficient ood adaptation in 3d inverse problems. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) ECCV 2024. pp. 432–455. Springer Nature Switzerland, Cham (2025)
5. Cohen, N., Manor, H., Bahat, Y., Michaeli, T.: From posterior sampling to meaningful diversity in image restoration. In: ICLR (2024)
6. Croce, F., Schlarmann, C., Singh, N.D., Hein, M.: Adversarially robust clip models can induce better (robust) perceptual metrics. In: IEEE SaTML (2025)

7. Desai, A.D., Schmidt, A.M., Rubin, E.B., Sandino, C.M., Black, M.S., Mazzoli, V., Stevens, K.J., Boutin, R., Re, C., Gold, G.E., et al.: Skm-tea: A dataset for accelerated mri reconstruction with dense image labels for quantitative clinical evaluation. In: *NeurIPS Datasets and Benchmarks Track* (2021)
8. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *IJCV* **88**, 303–338 (2010)
9. Ganz, R., Elad, M.: Clipag: Towards generator-free text-to-image generation. In: *WACV* (2024)
10. Harman, F., Selver, M.A., Baris, M.M., Canturk, A., Oksuz, I.: Deep learning-based meniscus tear detection from accelerated mri. *IEEE Access* **11**, 144349–144363 (2023)
11. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.I.: Robust compressed sensing mri with deep generative priors. *NeurIPS* (2021)
12. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdziel, M., Romero, A., Rabbat, M., Vincent, P., Pinkerton, J., Wang, D., Yakubova, N., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence* **2**(1), e190007 (2020), pMID: 32076662
13. Küstner, T., Hammernik, K., Rueckert, D., Hepp, T., Gatidis, S.: Predictive uncertainty in deep learning-based mr image reconstruction using deep ensembles: Evaluation on the fastmri data set. *Magnetic Resonance in Medicine* **92**(1), 289–302 (2024)
14. Li, Y., Mao, H., Girshick, R., He, K.: Exploring plain vision transformer backbones for object detection. In: *ECCV*. pp. 280–296. Springer (2022)
15. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: *CVPR* (2015)
16. Morshuis, J.N., Gatidis, S., Hein, M., Baumgartner, C.F.: Adversarial robustness of mr image reconstruction under realistic perturbations. In: Haq, N., Johnson, P., Maier, A., Qin, C., Würfl, T., Yoo, J. (eds.) *Machine Learning for Medical Image Reconstruction*. pp. 24–33. Springer International Publishing, Cham (2022)
17. Morshuis, J.N., Hein, M., Baumgartner, C.F.: Segmentation-guided mri reconstruction for meaningfully diverse reconstructions. In: Mukhopadhyay, A., Oksuz, I., Engelhardt, S., Mehro, D., Yuan, Y. (eds.) *Deep Generative Models*. pp. 180–190. Springer Nature Switzerland, Cham (2024)
18. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence* **39**(6), 1137–1149 (2016)
19. Santurkar, S., Ilyas, A., Tsipras, D., Engstrom, L., Tran, B., Madry, A.: Image synthesis with a single (robust) classifier. In: *NeurIPS* (2019)
20. Schlarmann, C., Singh, N.D., Croce, F., Hein, M.: Robust clip: Unsupervised adversarial fine-tuning of vision embeddings for robust large vision-language models. *ICML* (2024)
21. Solomon, E., Johnson, P.M., Tan, Z., Tibrewala, R., Lui, Y.W., Knoll, F., Moy, L., Kim, S.G., Heacock, L.: fastmri breast: A publicly available radial k-space dataset of breast dynamic contrast-enhanced mri. *Radiology: Artificial Intelligence* **7**(1), e240345 (2025)
22. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-end variational networks for accelerated mri reconstruction. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga,

- M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (eds.) MICCAI 2020. pp. 64–73. Springer International Publishing, Cham (2020)
23. Tibrewala, R., Dutt, T., Tong, A., Ginocchio, L., Lattanzi, R., Keerthivasan, M.B., Baete, S.H., Chopra, S., Lui, Y.W., Sodickson, D.K., et al.: Fastmri prostate: A public, biparametric mri dataset to advance machine learning for prostate cancer imaging. *Scientific Data* **11**(1), 404 (2024)
 24. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. In: *ICLR* (2019)
 25. Xie, Y., Li, Q.: Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. pp. 655–664. Springer Nature Switzerland, Cham (2022)
 26. Yen, C.Y., Singhal, R., Sharma, U., Ranganath, R., Chopra, S., Pinto, L.: Adaptive sampling of k-space in magnetic resonance for rapid pathology prediction. In: *ICML* (2024)
 27. Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y.W., Hansen, M.S., Lungren, M.P.: fastmri+, clinical pathology annotations for knee and brain fully sampled magnetic resonance imaging data. *Scientific Data* **9**(1), 152 (2022)