SIMPLE: Simultaneous Multi-plane Self-supervised Learning for Isotropic MRI Restoration from Anisotropic Data

Rotem Benisty[1 [0009−0005−0793−1072]], Yevgenia Shteynman[1 [0009−0000−4040−1982]], Moshe Porat[1 [0009−0007−1711−3828]], Anat Ilivitzki[2 [0000−0001−7770−4194]], and Moti Freiman[1 [0000−0003−1083−1548]]

[1] Technion – Israel Institute of Technology, Haifa, Israel
be.rotem@campus.technion.ac.il
[2] Rambam Health Care Campus, Haifa, Israel

Abstract. Magnetic resonance imaging (MRI) is vital for diagnosing abdominal and neurological conditions, yet conventional sequential slice acquisitions favor in-plane over through-plane resolution to minimize scan time and motion artifacts, leading to anisotropic data and reduced volumetric accuracy. Existing super-resolution (SR) techniques reconstruct isotropic images from anisotropic scans but often rely on simulated downsampling or limited 3D isotropic data, emphasizing through-plane interpolation rather than preserving full anatomy. We introduce SIMPLE, a Simultaneous Multi-Plane Self-Supervised Learning approach that directly restores isotropic MRI from real-world multi-plane acquisitions via adversarial training. Testing on OASIS-1 brain (n = 416) and Crohn's disease abdominal (n = 115) MRI datasets demonstrates SIMPLE's superiority in image fidelity and anatomical detail over state-of-the-art methods. Notably, SIMPLE achieved lower averaged Kernel Inception Distance (KID) scores than SMORE4 in both brain MRI (28.709 vs. 29.295) and abdominal MRI (17.435 vs. 20.724), retained higher-frequency details as confirmed by Fourier analysis, and was rated 1.5 points higher in the axial plane by radiologists. By improving volumetric analysis and 3D reconstructions, SIMPLE shows promise for enhancing diagnostic accuracy in pathologies demanding precise structural visualization. Our source code is publicly available at https://github.com/TechnionComputationalMRILab/SIMPLE.

Keywords: Isotropic MRI Restoration · Clinical MRI · Generative Adversarial Neural Networks .

1 Introduction

Magnetic resonance imaging (MRI) is central to medical diagnostics, offering high-resolution soft-tissue imaging across numerous conditions [6,9,18,10,7]. Yet achieving such detail often leads to anisotropic data, where differing in-plane and through-plane resolutions limit diagnostic precision and volumetric analysis

[5,21]. This stems from extensive 3D "k-space" sampling requirements: higher resolution prolongs scan times, making it infeasible in busy clinical settings [2]. Patient comfort, motion artifacts, and throughput demands further constrain acquisition, compelling a trade-off of higher in-plane resolution at the expense of lower through-plane resolution.

To compromise anisotropy and improve diagnostic accuracy, clinicians often acquire multiple high-resolution MRI scans in different planes [8]. While this multi-plane strategy offers a more comprehensive view, partial volume effects and reduced quality in the slice-selection direction still impair volumetric analysis and 3D reconstructions. Moreover, acquiring additional planes inevitably prolongs the overall scanning time.

Super-resolution (SR) has emerged as a promising solution by reconstructing isotropic, high-resolution images from anisotropic scans [12,17,19,25]. However, fully supervised SR requires isotropic ground-truth volumes, which are rarely feasible: body imaging suffers from respiratory motion, while neuroimaging faces long acquisition times.

Self-supervised SR methods eliminate the need for explicit HR ground truth by leveraging the existing anisotropic data – high-resolution (HR) in-plane but low-resolution (LR) through-plane – to form HR-LR pairs directly from the same scans. They generally fall into resampling-based and synthesis-based categories. Resampling-based approaches simulate even lower-resolution images from the acquired low-resolution (LR) data, then map them back to the original to predict HR [13]. Synthesis-based methods generate LR images from HR volumes and train on the resulting HR-LR pairs. For instance, Zhao et al. [27] downsampled the HR plane of anisotropic scans for their "SMORE" model, while Remedios et al. [20] introduced "SMORE4" by incorporating a learned point-spread function (PSF) and averaging two LR planes to form an isotropic volume. Similarly, slice-to-volume (SVR) [24] and patch-to-volume (PVR) [1] fuse axial, coronal, and sagittal slices into a 3D isotropic volume, and Liu et al. [16] further improved resolution by re-slicing and averaging multiple SR reconstructions. Despite their potential, most self-supervised SR methods rely on synthetic downsampling, small datasets, or indirect mappings, which can undermine performance. They also typically emphasize through-plane improvements, overlooking the broader 3D anatomy.

To address these gaps, we propose SIMPLE - a Simultaneous Multi-Plane Self-Supervised Learning approach that restores isotropic MRI directly from real-world anisotropic acquisitions without simulated downsampling or single-plane interpolation. SIMPLE's adversarial framework preserves anatomical detail across all three dimensions. Validated on large brain (OASIS-1) and abdominal (Crohn's disease) MRI datasets, SIMPLE consistently achieves superior image fidelity, high-frequency preservation, and expert radiologist preference compared to state-of-the-art methods. By supporting more accurate volumetric analysis and 3D reconstructions, SIMPLE has the potential to enhance diagnostics for abdominal and neurological pathologies alike.

2 Method

We present SIMPLE -a Simultaneous Multi-Plane Self-Supervised method for isotropic super-resolution of MRI volumes acquired with 2D protocols. This approach leverages an extended Generator–Discriminator framework, as illustrated in Figure 1. Specifically, we take an anisotropic volume $V_{\text{An-Iso}}$, apply linear interpolation to obtain a coarse isotropic volume V'_{Iso}, and subsequently generate a refined isotropic reconstruction $\hat{V}_{\text{Iso}}$.

During generator training, the output volume is sliced along the coronal, axial, and sagittal planes. Each plane is then processed by a dedicated discriminator, which uses real anisotropic data (with high in-plane resolution) to enforce realistic outputs for that plane. This strategy avoids the need for true 3D isotropic data or an approximate downsampling function. Simultaneously training multiple discriminators ensures consistent 3D restoration across all planes. Equation 1 formally defines the model, where GM is the generative module, L is the interpolation operator and i is the slice index within the isotropic volume:

$$\begin{aligned} &\hat{V}_{\text{Iso}} = GM(V'_{\text{Iso}}) = GM(L(V_{\text{An-Iso}})) \\ &\text{s.t:} \begin{cases} S_{\text{HR}_{\text{Cor}}} = \hat{V}_{\text{Iso}}[i,:,:] = HR(V'_{\text{Iso}}[i,:,:]) \\ S_{\text{HR}_{\text{Ax}}} = \hat{V}_{\text{Iso}}[:,i,:] = HR(V'_{\text{Iso}}[:,i,:]) \\ S_{\text{HR}_{\text{Sag}}} = \hat{V}_{\text{Iso}}[:,:,i] = HR(V'_{\text{Iso}}[:,:,i]) \end{cases} \end{aligned} \tag{1}$$

2.1 System Architecture

Pre-Processing We first linearly interpolate the anisotropic volume to produce a coarse isotropic volume, though large slice spacing degrades voxel accuracy. To handle the high memory demands of processing the entire volume, we then extract overlapping 3D patches of size $64 \times 64 \times 64$ for subsequent model training.

Networks Architecture We adopt a 3D generator based on 3D U-Net [4] to transform $64 \times 64 \times 64$ single-channel, low-quality patches into high-quality counterparts. The number of discriminators depends on available planes (three for coronal-axial-sagittal or two for coronal-axial), and each operates on 2D slices extracted from the generator's 3D output. Following a conditional patch-discriminator design [11], both real and synthetic samples contain two channels: the low-resolution input slice and either the generator's output (synthetic) or an authentic high-resolution slice (real).

However, our data consist of anisotropic acquisitions obtained at different times, often with motion artifacts and mismatched anatomical coverage. Consequently, the additional planes cannot directly serve as real high-resolution references for the discriminator. To address this, we pre-train a 2D single-image super-resolution model (ATME [23]) on the anisotropic acquisitions for each plane and evaluate it on the linearly interpolated low-resolution volume to generate 'real' high-resolution slices in each plane. These synthetic-yet-consistent

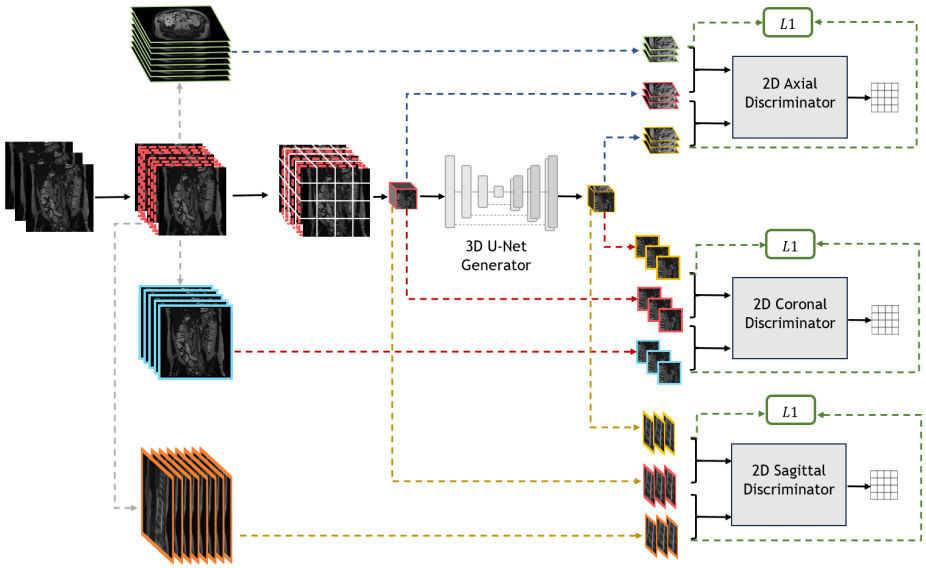

Fig. 1: Model Architecture. The framework incorporates a 3D generator and three (or two, depending on plane availability) discriminators operating in parallel. The dashed red volume indicates the linearly interpolated input, while the blue, green, and orange frames depict high-quality coronal, axial, and sagittal slices generated by ATME. Red, blue, and yellow dashed arrows indicate sampling along the coronal, axial, and sagittal planes, respectively. Gray dashed arrows show the ATME flow and green dashed arrows highlight the connections to consistency loss.

references replace the original scans in the conditional discriminator, ensuring matched real–synthetic pairs while circumventing misalignment from multi-plane acquisitions.

Loss Functions During training, we employ separate objectives for the generator and discriminators to balance realism and consistency across all planes. The generator's total loss equally weights coronal, axial, and sagittal components, each consisting of an adversarial term $\mathcal{L}_{G_{\text{ADV}}}$ (from the discriminator's classification matrix) and a consistency term $\mathcal{L}_{L_1}$ (the L1 distance). The discriminators rely solely on the adversarial loss, implemented via a least squares generative adversarial network (LSGAN). Equations 2 and 3 detail the generator and discriminator objectives, with x as the input and y as the target.

$$\mathcal{L}_{\text{G}_{\text{ADV}}} = \mathcal{L}_{\text{LSGAN}}(G, D) = \mathbb{E}_x \left[\left(D(x, G(x)) - 1 \right)^2 \right] \tag{2}$$

$$\mathcal{L}_{\text{D}_{\text{ADV}}} = \mathcal{L}_{\text{LSGAN}}(D) = \frac{1}{2}\mathbb{E}_{x,y} \left[\left(D(x, y) - 1 \right)^2 \right] + \frac{1}{2}\mathbb{E}_x \left[\left(D(x, G(x)) \right)^2 \right] \tag{3}$$

Equation 4 represents our total generator loss:

$$\mathcal{L}_{\mathrm{G}} = \alpha \left[\mathcal{L}_{\mathrm{G}_{\mathrm{ADV}_{\mathrm{cor}}}} + \lambda \mathcal{L}_{L_{1_{\mathrm{cor}}}}\right] + \beta \left[\mathcal{L}_{\mathrm{G}_{\mathrm{ADV}_{\mathrm{ax}}}} + \lambda \mathcal{L}_{L_{1_{\mathrm{ax}}}}\right] + \gamma \left[\mathcal{L}_{\mathrm{G}_{\mathrm{ADV}_{\mathrm{sag}}}} + \lambda \mathcal{L}_{L_{1_{\mathrm{sag}}}}\right] \tag{4}$$

where α, β, γ denote the weights for each plane component and are empirically set to 0.5.

3 Experiments

3.1 Datasets

Abdomen Dataset. After Institutional Review Board (IRB) approval, we gathered 115 consecutive 2D MRI scans of Crohn's Disease (CD) patients from a local hospital, acquired on a GE 3T scanner. Coronal and axial FIESTA sequences were used, each with 5 mm slice spacing and thickness. The average pixel spacing is 0.78 mm (coronal) and 0.76 mm (axial).
OASIS-1 Dataset. We also used the publicly available OASIS-1 dataset of 416 brain MRIs from adults with varying cognitive status. Each T1-weighted sagittal scan features 1.25 mm slice spacing and 1 mm in-plane resolution, plus a co-registered 3D volume resampled to 1 mm isotropic. We generated anisotropic 2D MRI sequences for each plane by setting a 3 mm slice spacing and preserving the 1 mm in-plane resolution.

3.2 Experimental Methodology

We compare SIMPLE against three baselines: linear interpolation, SMORE4 [20], and a method based on Liu et al. [16] that applies ATME-based super-resolution independently to each plane before averaging the resulting volumes. To ensure a fair and clinically relevant comparison, we evaluate results on the coronal plane for the abdomen dataset and the sagittal plane for the brain dataset.

To assess qualitative performance, we analyze multi-view slices in the coronal, axial, and sagittal planes, providing a direct visualization of each method's reconstruction. Since the abdomen dataset lacks true 3D isotropic volumes, no ground-truth slices are available for direct reference. To further evaluate volumetric consistency, we generate straight multi-planar reconstructions (MPRs) of the small intestine using radiologist-annotated Terminal Ileum centerlines [14]. Additionally, we apply 2D Fourier transforms to slices from planes other than the primary evaluation plane to assess high-frequency detail preservation, such as texture and edge sharpness. To incorporate clinical validation, a senior radiologist, blinded to the methods, scores slices from 10 abdomen cases across coronal, axial, and sagittal views using a 1–5 Likert scale, where 1 represents non-diagnostic quality and 5 indicates excellent visualization.

For quantitative evaluation, we employ distribution-based metrics, as PSNR and SSIM have been shown to be unreliable for GAN-based medical imaging [26],

and the absence of ground-truth references in the abdomen dataset limits the applicability of full-reference metrics. We compute Fréchet Inception Distance (FID) [15,22], Kernel Inception Distance (KID), and the L1 distance between Inception Scores (IS) [3]. For all metrics, lower scores correspond to better performance. For the abdomen dataset, scores are computed separately for coronal and axial slices and then averaged. For the brain dataset, KID scores are calculated across 2D slices in each plane and averaged across cases to provide a 3D evaluation.

4 Results

Multi View Slices: Figure 2 compares isotropic abdominal MRI volumes reconstructed by all methods across the three primary planes. Interpolation and ATME enhance image quality primarily in the plane of their original anisotropic acquisition but provide limited improvement in the perpendicular planes. SIMPLE, SMORE4, and averaged ATMEs demonstrate better resolution across all planes, yet SIMPLE achieves the most consistent enhancement, particularly in planes perpendicular to the acquisition plane, where it sharpens edges and improves contrast more effectively.

In the acquisition plane, SIMPLE and SMORE4 yield similar results, though minor differences in organ location are observed, with SMORE4 being more sensitive to organ shadows and noise. While SIMPLE has no explicit constraints in the sagittal plane for the abdomen dataset, it still produces improved resolution. For the brain dataset, the differences between SIMPLE and other methods are less pronounced than in the abdomen dataset due to the smaller slice spacing and minimal organ motion, suggesting that SIMPLE's advantages are most evident in cases with greater through-plane anisotropy.

Fourier Domain: Figure 3 presents the Fourier domain representation of axial and coronal slices from isotropic volumes reconstructed by all methods. SIMPLE's Fourier transform closely matches that of natural images, demonstrating a well-balanced distribution of high and low frequencies. In contrast, interpolation introduces visible horizontal stripes in the axial abdominal slices, indicative of retained anisotropic resolution. Similarly, SMORE4 and averaged ATMEs exhibit a prominent horizontal stripe near the center, suggesting residual artifacts from the original anisotropic acquisition.

All methods display vertical and horizontal lines at the center, likely resulting from zero-padding effects. In the brain dataset, the coronal and axial slices generated by competing methods show a noticeable reduction in high-frequency components, deviating further from the original Fourier domain. These results indicate that SIMPLE better preserves fine structural details, particularly in anisotropic cases, reinforcing its ability to generate more natural and high-fidelity isotropic reconstructions.

Straight Multi Planar Reconstruction (MPR): Figure 4 presents straight MPR reconstructions from five volumes: the original anisotropic acquisition and isotropic reconstructions generated by all methods. SIMPLE produces the

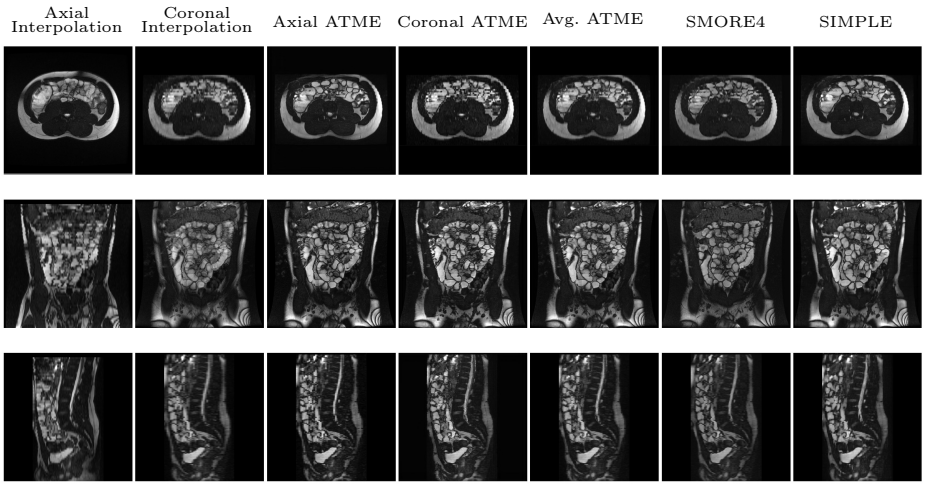

Fig. 2: Multi-plane comparison of isotropic slices generated by SIMPLE versus five baselines: (1) linear interpolation from an anisotropic axial volume, (2) linear interpolation from an anisotropic coronal volume, (3) ATME trained on the axial volume, (4) ATME trained on the coronal volume, (5) averaged ATME across both planes, (6) SMORE4, and (7) SIMPLE. All methods except (1) were applied to an anisotropic coronal volume. SIMPLE yields sharper details and more consistent 3D structures, highlighting its advantage over existing methods.

smoothest and sharpest reconstructions, reducing discontinuities that arise in both the anisotropic volume and competing isotropic methods. These discontinuities, observed in SMORE4 and averaged ATMEs, stem from larger inter-slice gaps and insufficient through-plane reconstruction. SIMPLE's ability to generate more coherent and continuous structures highlights its effectiveness in reducing slice-wise inconsistencies and better preserving anatomical integrity in 3D volumetric reconstructions.

Likert Scale: Table 2 displays the Likert scale rankings across all planes. Linear interpolation performed the worst, with an average score of approximately 1. SIMPLE and SMORE4 achieved similarly high scores in the coronal plane and lower scores in the sagittal plane, though SIMPLE ranked slightly higher in both. In the axial plane, SIMPLE outperformed SMORE4 by 1.5 points on average, indicating better anatomical alignment and fewer artifacts, essential for diagnostic reliability.

Distribution-Based Metrics: Tables 3 and 1 summarize the metric scores for each method on the brain and abdomen datasets. In the abdomen dataset, SIMPLE scores lower than interpolation in both planes and achieves lower scores than SMORE4 and averaged ATMEs in the axial plane, resulting in the lowest overall average scores among all methods. By contrast, SMORE4 and averaged ATMEs perform well in the coronal plane but degrade more significantly in

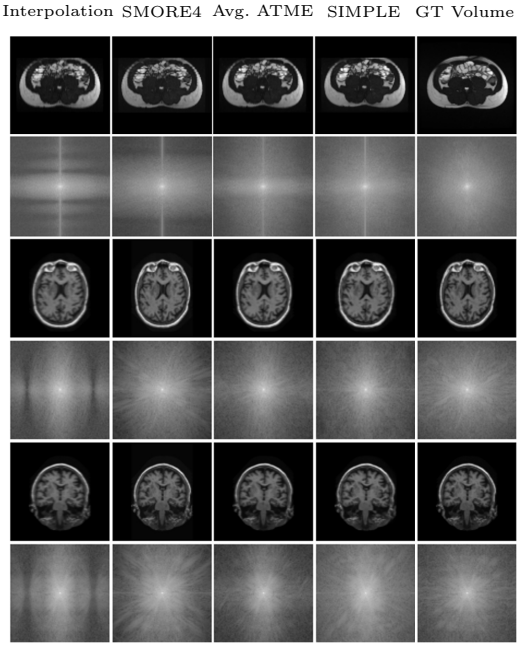

Fig. 3: Fourier representation of axial and coronal slices reconstructed by SIMPLE and other methods from anisotropic volumes, compared to a high-resolution reference. SIMPLE preserves high-frequency details, indicating close fidelity to real images.

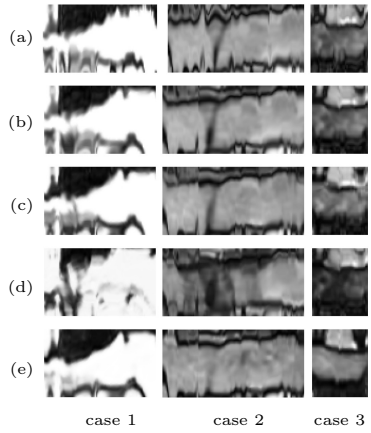

Fig. 4: Straight MPR of the terminal ileum in Crohn's disease patients from coronal volumes: (a) anisotropic volume, (b) linear interpolation, (c) averaged ATMEs, (d) SMORE4, and (e) SIMPLE. SIMPLE provides smoother, sharper reconstructions, illustrating its potential for improved clinical assessment of Crohn's disease.

Table 1: Distribution-based metrics for the abdomen dataset, listing plane-specific results (coronal, axial) and the overall average across both planes.

	Interpolation		**SMORE4**		**Avg. ATME**		**SIMPLE**	
	Coronal	*Axial*	*Coronal*	*Axial*	*Coronal*	*Axial*	*Coronal*	*Axial*
KID	3.829	48.185	0.894	40.554	0.639	37.213	2.383	**32.487**
	26.007		20.724		18.926		**17.435**	
IS	0.031	0.023	0.001	0.033	0.005	0.022	0.003	**0.006**
	0.027		0.017		0.013		**0.004**	
FID	17.292	28.964	14.897	23.938	15.387	22.969	16.834	**21.316**
	23.128		19.417		19.178		**19.075**	

Table 2: Average and standard deviation of Likert scale rankings for the three main planes across competing methods.

	Interpolation	**SMORE4**	**SIMPLE**
Axial	1	2.182 ± 1.168	**3.727 ± 0.786**
Coronal	1.231 ± 0.438	4.231 ± 0.438	**4.615 ± 0.65**
Sagittal	1	1.5 ± 0.527	**2.2 ± 0.789**

the axial plane, leading to greater inter-plane variability. For the brain dataset, SIMPLE consistently yields significantly lower KID scores than interpolation in both 2D and 3D and slightly outperforms the other methods, especially in the coronal and axial planes.

Table 3: KID scores for the brain dataset, showing the average KID per plane and the overall mean across all planes, aggregated by case.

	Interpolation			**SMORE4**			**Avg. ATME**			**SIMPLE**		
KID	*Coronal*	*Axial*	*Sagittal*	*Coronal*	*Axial*	*Sagittal*	*Coronal*	*Axial*	*Sagittal*	*Coronal*	*Axial*	*Sagittal*
2D	0.89	0.964	0.386	0.47	0.495	0.354	0.418	0.455	0.241	**0.396**	**0.422**	**0.231**
3D	30.425			29.295			29.36			**28.709**		

5 Conclusion

We presented SIMPLE, a self-supervised simultaneous multi-plane learning approach for isotropic MRI reconstruction from anisotropic clinical acquisitions. By leveraging real-world multi-plane data and adversarial training, SIMPLE preserves full anatomical structure without relying on simulated downsampling. Extensive experiments on brain and abdominal MRI datasets demonstrate its superiority over state-of-the-art methods in image fidelity, high-frequency detail preservation, and radiologist evaluations. These results highlight SIMPLE's potential to enhance volumetric analysis and improve diagnostic accuracy for abdominal and neurological pathologies, making it a promising solution for high-resolution MRI reconstruction in clinical practice.

Acknowledgments. This study was supported by Kamin Grant No. 73249 from the Israel Innovation Authority.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alansary, A., Rajchl, M., McDonagh, S.G., Murgasova, M., Damodaram, M., Lloyd, D.F., Davidson, A., Rutherford, M., Hajnal, J.V., Rueckert, D., et al.: Pvr: patch-to-volume reconstruction for large area motion correction of fetal mri. IEEE transactions on medical imaging **36**(10), 2031–2044 (2017)
2. Bankman, I.: Handbook of medical image processing and analysis. Elsevier (2008)
3. Barratt, S., Sharma, R.: A note on the inception score. arXiv preprint arXiv:1801.01973 (2018)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19. pp. 424–432. Springer (2016)
5. Dev, H., Zhu, C., Sharbatdaran, A., Raza, S.I., Wang, S.J., Romano, D.J., Goel, A., Teichman, K., Moghadam, M.C., Shih, G., et al.: Effect of averaging measurements from multiple mri pulse sequences on kidney volume reproducibility in autosomal dominant polycystic kidney disease. Journal of Magnetic Resonance Imaging **58**(4), 1153–1160 (2023)

6. Donato, H., França, M., Candelária, I., Caseiro-Alves, F.: Liver mri: from basic protocol to advanced techniques. European journal of radiology **93**, 30–39 (2017)
7. Francis, S.T., Selby, N.M., Taal, M.W.: Magnetic resonance imaging to evaluate kidney structure, function, and pathology: Moving toward clinical application. American Journal of Kidney Diseases **82**(4), 491–504 (2023)
8. Gore, R.M., Levine, M.S.: Textbook of gastrointestinal radiology. Elsevier Health Sciences (2021)
9. Gourtsoyiannis, N.C., Aschoff, P.: Clinical MRI of the abdomen: why, how, when. Springer (2011)
10. Harrington, K.A., Shukla-Dave, A., Paudyal, R., Do, R.K.: Mri of the pancreas. Journal of Magnetic Resonance Imaging **53**(2), 347–359 (2021)
11. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
12. Jia, Y., Gholipour, A., He, Z., Warfield, S.K.: A new sparse representation framework for reconstruction of an isotropic high spatial resolution mr volume from orthogonal anisotropic resolution scans. IEEE transactions on medical imaging **36**(5), 1182–1193 (2017)
13. Jog, A., Carass, A., Prince, J.L.: Self super-resolution for magnetic resonance images. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part III 19. pp. 553–560. Springer (2016)
14. Kanitsar, A., Fleischmann, D., Wegenkittl, R., Felkel, P., Groller, E.: CPR-curved planar reformation. IEEE (2002)
15. Kastryulin, S., Zakirov, J., Pezzotti, N., Dylov, D.V.: Image quality assessment for magnetic resonance imaging. IEEE Access **11**, 14154–14168 (2023)
16. Liu, Y., Liu, Y., Vanguri, R., Litwiller, D., Liu, M., Hsu, H.Y., Ha, R., Shaish, H., Jambawalikar, S.: 3d isotropic super-resolution prostate mri using generative adversarial networks and unpaired multiplane slices. Journal of Digital Imaging **34**, 1199–1208 (2021)
17. Masutani, E.M., Bahrami, N., Hsiao, A.: Deep learning single-frame and multiframe super-resolution for cardiac mri. Radiology **295**(3), 552–561 (2020)
18. Nowak, S., Mesropyan, N., Faron, A., Block, W., Reuter, M., Attenberger, U.I., Luetkens, J.A., Sprinkart, A.M.: Detection of liver cirrhosis in standard t2-weighted mri using deep transfer learning. European radiology **31**(11), 8807–8815 (2021)
19. Peng, C., Zhou, S.K., Chellappa, R.: Da-vsr: domain adaptable volumetric super-resolution for medical images. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24. pp. 75–85. Springer (2021)
20. Remedios, S.W., Han, S., Zuo, L., Carass, A., Pham, D.L., Prince, J.L., Dewey, B.E.: Self-supervised super-resolution for anisotropic mr images with and without slice gap. In: International Workshop on Simulation and Synthesis in Medical Imaging. pp. 118–128. Springer (2023)
21. Ristow, O., Steinbach, L., Sabo, G., Krug, R., Huber, M., Rauscher, I., Ma, B., Link, T.M.: Isotropic 3d fast spin-echo imaging versus standard 2d imaging at 3.0 t of the knee—image quality and diagnostic performance. European radiology **19**, 1263–1272 (2009)
22. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)

23. Solano-Carrillo, E., Rodriguez, A.B., Carrillo-Perez, B., Steiniger, Y., Stoppe, J.: Look atme: the discriminator mean entropy needs attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 787–796 (2023)
24. Uus, A., Zhang, T., Jackson, L.H., Roberts, T.A., Rutherford, M.A., Hajnal, J.V., Deprez, M.: Deformable slice-to-volume registration for motion correction of fetal body and placenta mri. IEEE transactions on medical imaging **39**(9), 2750–2759 (2020)
25. Wang, X., Song, Z., Zhu, Y., Wang, S., Zhang, L., Shen, D., Wang, Q.: Inter-slice super-resolution of magnetic resonance images by pre-training and self-supervised fine-tuning (2024), https://arxiv.org/abs/2406.05974
26. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: A review. Medical image analysis **58**, 101552 (2019)
27. Zhao, C., Dewey, B.E., Pham, D.L., Calabresi, P.A., Reich, D.S., Prince, J.L.: Smore: a self-supervised anti-aliasing and super-resolution algorithm for mri using deep learning. IEEE transactions on medical imaging **40**(3), 805–817 (2020)