Abstract #06430 ISMRM ABSTRACT

Enhancement of Mid-Field (0.6T) T2-Weighted Prostate Scans Using a Two-Stage Refinement Framework

Primary: Analysis Methods - Image Enhancement 〉 **Secondary:** Physics & Engineering - Low-Field MRI 〉 **PowerPitch Oral** · 2 min + **Digital Poster** · 60 min | Next-Generation MRI Image Enhancement · Tuesday, May 12 at 04:00 PM 〉 **Keywords:** IMAGE QUALITY ENHANCEMENT 0.6T TWO-STAGE REFINEMENT FLOW MATCHING MID-FIELD PROSTATE MRI

Donghang Lyu [1], **Martijn Nagtegaal**[2], **Ece Ercan**[3], **Zhong Li**[4], **Mart WJ van Straten**[2,5], **Marius Staring**[1], **Andrew Webb**[2], **Matthias van Osch**[2], **Peter Börnert**[2,6], **Yiming Dong**[2]

[1]Department of Radiology, Leiden University Medical Center, Leiden, Netherlands
[2]Department of Radiology, C.J. Gorter MRI Center, Leiden University Medical Center, Netherlands
[3]Philips Healthcare, Best, Netherlands
[4]Great Bay University, Dongguan, China
[5]Department of Biomedical Engineering, Eindhoven University of Technology, Eindhoven, Netherlands
[6]Philips Innovative Technologies, Hamburg, Germany

Presenting Author: *Donghang Lyu (d.lyu@lumc.nl)*

Impact

By applying the proposed two-stage refinement framework, mid-field prostate MRI with heavy noise and low resolution can be effectively enhanced, achieving image quality comparable to higher-field MRI.

Synopsis

Motivation: Mid-field MRI shows potential for prostate imaging. Deep learning–based denoising and super-resolution methods can effectively enhance its image quality.

Goals: To enhance 0.6T T2-weighted prostate images to a quality comparable to 3T data.

Approach: A two-stage refinement framework trained on simulated 0.6T data is proposed, enabling more effective detail recovery compared to single-stage image enhancement.

Results: By incorporating a second-stage refinement network, both denoising and super-resolution are effectively achieved, resulting in enhanced image quality comparable to that of higher-field MRI.

Introduction

High-quality T2-weighted (T2w) imaging is essential for prostate anatomy assessment and lesion characterization, typically performed at 3T using turbo spin echo (TSE) sequences. Midfield MRI (e.g., 0.6T) is attractive for prostate imaging due to the improved safety and reduced magnetic susceptibility artifacts. However, its inherently lower SNR limits the achievable spatial resolution and image quality. Acquiring data at lower in-plane resolution can improve SNR and shorten scan time, but may compromise the fine structural details required for clinical diagnosis.

Deep learning–based image enhancement methods effectively denoise and remove undersampling artifacts but often lead to over-smoothing [5]. When super-resolution is incorporated, it becomes challenging to train a single network that simultaneously achieves effective noise suppression and fine-structure preservation. To address this, we propose a two-stage quality enhancement framework for 0.6T prostate T2w MRI to fulfill clinical needs. In the first stage, a U-Net [1] model performs denoising and super-resolution on the noisy, lower-resolution input. In the second stage, a flow matching [2] pipeline refines the image by enhancing fine structural details and further improving resolution, leveraging its generative modeling and distribution alignment capabilities.

Method

The FastMRI prostate dataset [3] was used to train the proposed two-stage framework, assuming minimal field-dependent T2 values difference between 0.6T and 3T, and inference was performed on in-house 0.6T data. To simulate 0.6T data quality from 3T T2w prostate data, several preprocessing steps were applied. Specifically, the FastMRI T2w prostate images (DICOM format) with in-plane resolution 0.56 mm^2 were downsampled to the resolution of the acquired mid-field scans by cropping in k-space, followed by the addition of synthetic noise corresponding to the SNR of the acquired mid-field scans at each resolution. Zero-padding was applied to the k-space data to generate the training inputs, matching the matrix size between the noisy low-resolution training inputs and the target 3T high-resolution high SNR data.

The overall refinement framework consisted of two stages, as illustrated in Fig.2. In the first stage, we employed a 4-level U-Net architecture for denoising and super-resolution. To further improve representational capacity, a multi-head self-attention block [4] was incorporated at the bottleneck to expand the network's receptive field. The second stage was introduced to recover missing details and further enhance image resolution. Unlike commonly-used diffusion models, a flow matching model [2] enabled fast, single-pass generation without stochastic noise during sampling, resulting in more efficient inference. Accordingly, we employed it to learn a continuous transformation from the first-stage output to the high-resolution ground truth by modeling probability flow dynamics, allowing efficient and deterministic refinement that restores fine structural details without stochastic sampling.

The training FastMRI prostate dataset [3] was acquired on two clinical 3T systems (MAGNETOM Vida, Siemens, Germany) from 312 male patients. Axial TSE sequences were used for acquisition. For validation, three healthy volunteers were included in this study and scanned on a prototype 0.6T BlueSeal Philips scanner (Best, the Netherlands) equipped with 45mT/m (slew max 200 mT/m/ms) gradients and body receive arrays (posterior and anterior, 16 elements each); one of them was scanned with resolution 0.56×0.80 mm^2 in-plane, which has lower SNR and two of them were scanned with 1.12×1.12 mm^2 with higher SNR. Full imaging parameters are summarized in Fig.1. During training, both stages were optimized jointly, using a weighted sum of L1 and SSIM losses for the first stage and an MSE loss for the second stage.

Results

Fig.2 presents example qualitative results obtained during training (simulated input resolution 1.12mm^2). The output from the first stage demonstrates effective denoising; however, the capability of the single network was limited, and some fine anatomical details were missing, leading to an overall over-smoothed appearance, as shown in the lower-left example. In the second stage, owing to the flow-matching property of learning the velocity field between two distributions, fine structures are further refined after inference. Fig.3 and Fig.4 show results of applying the pre-trained two-stage framework to in-house 0.6T data with two different input resolutions, highlighting the performance for super-resolution and denoising tasks, respectively.

Discussion and conclusion

Overall, the proposed two-stage refinement method demonstrated strong performance in effective denoising and detail restoration. Since a single model often struggles to perform both denoising and super-resolution simultaneously, introducing an additional refinement stage provides an effective solution. Furthermore, the use of simulated data was proven to be both feasible and beneficial, as supported by previous studies [6,7] and confirmed by our results. Although the proposed two-stage refinement framework demonstrated promising image quality enhancement, future work could be extended by incorporating reconstruction from raw k-space data.

Acknowledgements

This collaboration project (specifically MN and YD) is co-funded by the PPS Allowance made available by Health~Holland, Top Sector Life Sciences & Health, to stimulate public-private partnerships. This project is a collaboration between Philips and LUMC.

References

1. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation. International Conference on Medical Image Computing and Computer-Assisted Intervention. Berlin: Springer; 2015. p 234–241.

2. Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., & Le, M. Flow Matching for Generative Modeling. In The Eleventh International Conference on Learning Representations.

3. Tibrewala, R., Dutt, T., Tong, A., Ginocchio, L., Lattanzi, R., Keerthivasan, M. B., ... & Johnson, P. M. (2024). FastMRI Prostate: A public, biparametric MRI dataset to advance machine learning for prostate cancer imaging. Scientific data, 11(1), 404.

4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30.

5. Chen, Z., Pawar, K., Ekanayake, M., Pain, C., Zhong, S., & Egan, G. F. (2023). Deep learning for image enhancement and correction in magnetic resonance imaging—state-of-the-art and challenges. Journal of Digital Imaging, 36(1), 204-230.

6. Ma, Q., Meng, Q., Qiao, M., Matthews, P. M., O'Regan, D. P., & Bai, W. (2025, September). CardiacFlow: 3D+ t Four-Chamber Cardiac Shape Completion and Generation via Flow Matching. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 89-99). Cham: Springer Nature Switzerland.

7. Ayaz, A., Boonstoppel, R., Lorenz, C., Weese, J., Pluim, J., & Breeuwer, M. (2024). Effective deep-learning brain MRI super resolution using simulated training data. Computers in Biology and Medicine, 183, 109301.

Figures and Tables

	TR (s)	TE (ms)	In-plane resolution (mm^2)	Slice thickness (mm)	FOV (mm^2)	Slices/volume
Training dataset (3T)	3.5–7.2	100	0.56 × 0.56	3	180 × 180	30–36
Validation dataset (0.6T)	3.2	120	0.56 × 0.80	3	200 × 278	21
	3.1	120	1.12 × 1.12	3	200 × 281	25

Scan for high-resolution version

Figure 1: *Scan parameters of the 3T FastMRI training set and 0.6T validation set.*

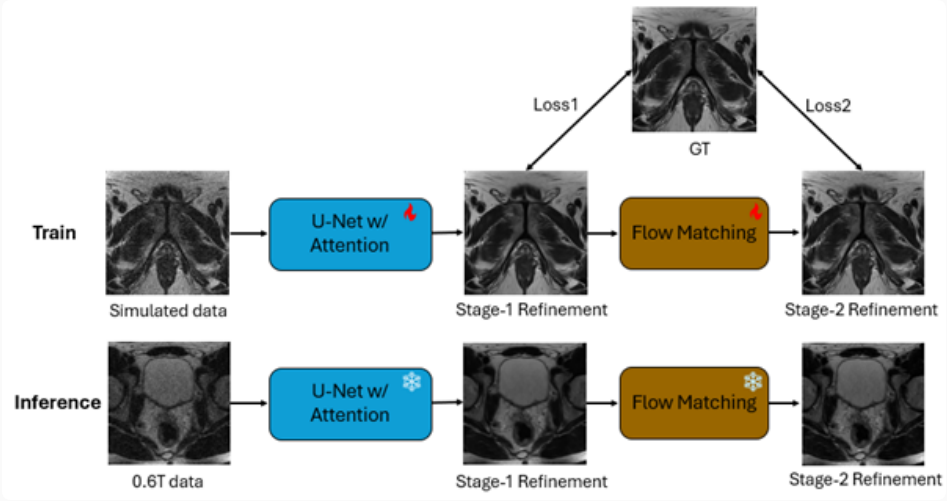

Scan for high-resolution version

Figure 2: *Overall training and inference pipeline of the two-stage refinement framework. The upper part shows training with simulated low-resolution and noisy FastMRI samples, where each stage's output is compared with the high-resolution ground truth (GT) to compute loss. The lower part shows inference on 0.6T cases using frozen networks. Fire and snowflake icons indicate trainable and frozen components, respectively.*

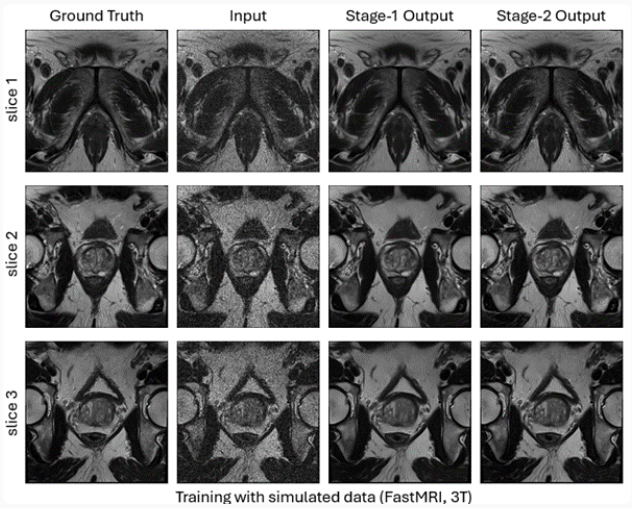

Scan for high-resolution version

Figure 3: *Qualitative results of the two-stage refinement during training. The simulated 0.6T input (1.12×1.12 mm^2, 2nd column) exhibits lower resolution and higher noise compared to the ground truth (1st column). Stage 1 (3rd column) effectively denoises but appears over-smoothed, while Stage 2 (4th column) further enhances the resolution and restores structural details, closely resembling the ground truth.*

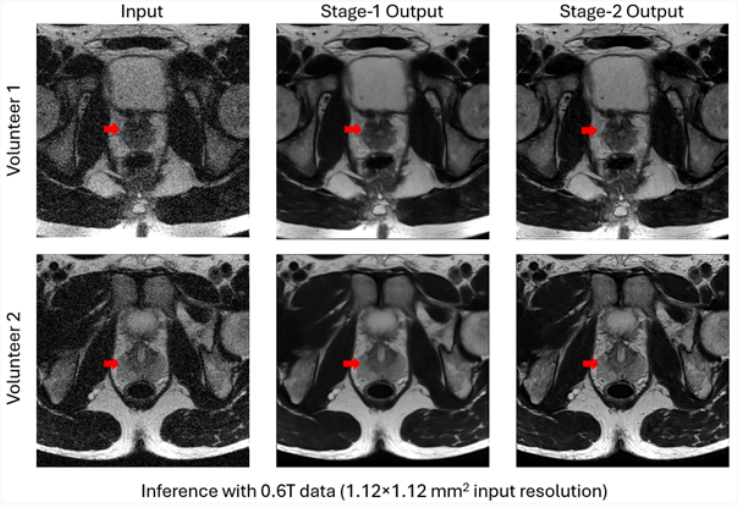

Scan for high-resolution version

Figure 4: *Qualitative results of the two-stage refinement during inference (input 1.12×1.12 mm^2). The input consists of 0.6T in-house data, and the overall refinement exhibits the same trend as observed during training: Stage 1 effectively performs denoising, while Stage 2 further enhances resolution and restores fine details (red arrows). These results also demonstrate the feasibility of using simulated data for training and doing inference on the real data.*

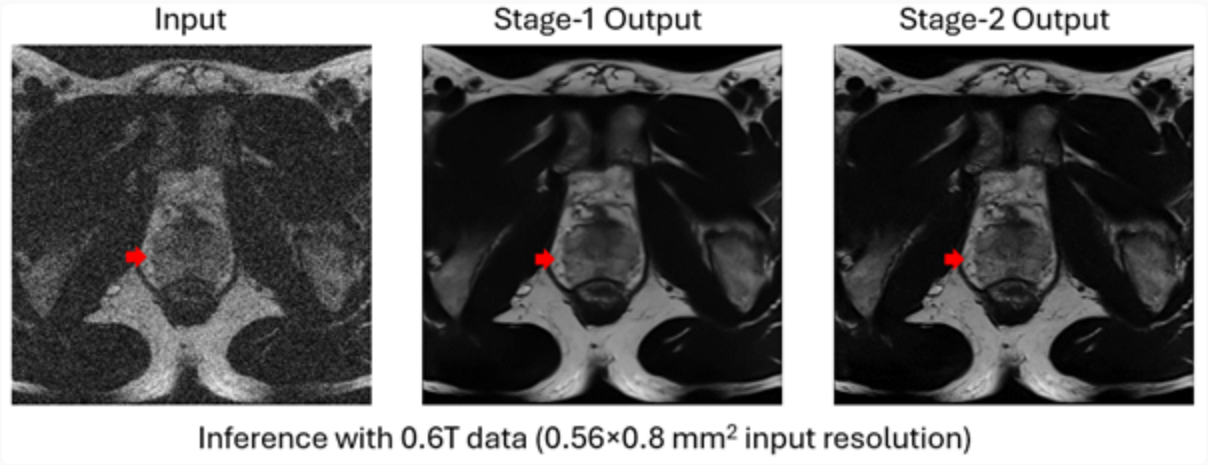

Scan for high-resolution version

Figure 5: *Qualitative results of the two-stage refinement during inference (input resolution: 0.56×0.80 mm²). As the in-plane resolution increased in the 0.6T scan compared to figure 4, the SNR decreased; however, the proposed framework effectively denoised the image in Stage 1 and restored fine structural details in Stage 2. (red arrows).*