Contrast-Invariant Reference-Based Slice Thickness Estimation in MRI via Spectral Matching

Samuel W. Remedios[1], Shuwen Wei[1], Aaron Carass[1], Jerry L. Prince[1], and Blake E. Dewey[2]

[1] The Image Analysis and Communications Laboratory, Johns Hopkins University
samuel.remedios@jhu.edu
[2] Department of Neurology, Johns Hopkins Hospital

Abstract. Accurate modeling of the slice selection profile is essential for super-resolution and resolution quantification in 2D-acquired MRI, where slice thickness and spacing are independent. Assuming that the slice spacing is equal to the slice thickness leads to systematic errors, yet estimating slice thickness from reconstructed images is ill-posed. We propose PRISM (Profile Recovery via Invariant Spectral Matching), a contrast-invariant method that estimates slice thickness by matching whole-volume spectral characteristics between a low-resolution acquisition and a degraded high-resolution reference. The reference image need not be the same contrast or subject and does not need registration nor voxel-wise correspondence to the acquired image. PRISM applies the rank transform to suppress contrast differences and compares the average 1D log power spectra along the slice axis. The thickness parameter is recovered via a 1D search without gradient-based optimization. Across controlled simulations, ablations, downstream super-resolution experiments, and real 2D-acquired data, PRISM accurately estimates slice thickness (within 0.09 mm for 3 mm data and within 0.64 mm for 10 mm data) and outperforms existing kernel estimation approaches, yielding reconstruction performance statistically indistinguishable from using the true slice profile. Code is available at https://github.com/sremedios/PRISM.

Keywords: MRI · slice thickness · slice profile · blur kernel

1 Introduction

Magnetic resonance imaging (MRI) is often acquired using 2D multi-slice protocols due to favorable signal-to-noise efficiency and practical scan-time constraints [1, 2, 5, 19]. In these acquisitions, slice thickness and slice spacing are independent design parameters [4, 17]. The through-plane resolution is defined by the slice thickness, and is therefore not necessarily equal to the slice spacing. The slice thickness is determined by the full-width-half-max (FWHM) of the point spread function (alternatively, the "kernel") induced by the slice selection profile. Inaccurate modeling of this profile biases resolution quantification and degrades downstream tasks such as super-resolution (SR).

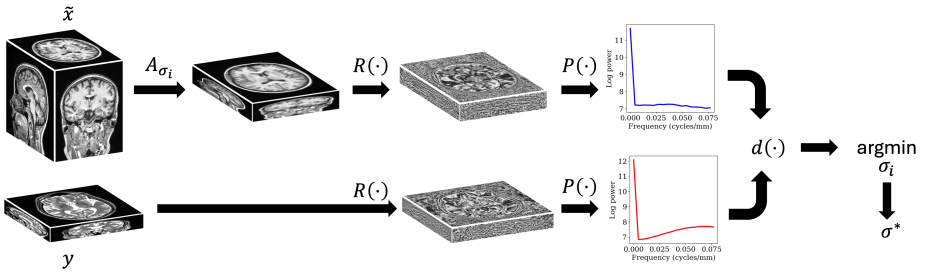

Fig. 1. PRISM uses the HR reference image $\tilde{x}$ and the LR image y to produce the standard deviation σ^* of a Gaussian kernel that achieves the minimal error d over a set of σ_i. The rank transform R is used as a contrast-invariant resolution-sensitive feature extractor, and the log spectral power P is used to characterize resolution.

Estimating slice thickness directly from reconstructed images is challenging and ill-posed. Blind kernel estimation methods are prone to degenerate solutions [6, 13], while existing approaches often fix degradation heuristically, learn implicit kernels, or rely on correspondence-based matching that is not applicable across subjects or contrasts in MRI [3, 8, 12, 14, 15, 24]. Consequently, explicit and interpretable slice thickness estimation in MRI remains underexplored, with only limited prior work [9].

Slice selection acts as a 1D low-pass filter along the slice axis, attenuating high frequencies. While anatomy and contrast affect spatial organization and intensity distributions, blur produces a predictable spectral decay. This suggests that through-plane resolution can be characterized using whole-volume spectral statistics without requiring anatomical correspondence.

Based on this insight, we propose PRISM (Profile Recovery via Invariant Spectral Matching), a contrast-invariant framework for slice thickness estimation. Given a low-resolution (LR) acquisition and a high-resolution (HR) reference volume (not necessarily aligned or from the same subject), PRISM generates candidate Gaussian slice profiles and estimates slice thickness by matching 1D log power spectra of rank-transformed volumes along the slice axis (Fig. 1). The optimal thickness is obtained via a 1D exhaustive search.

We validate PRISM through simulations, ablation studies, downstream super-resolution experiments, and real 2D-acquired LR volumes. Across datasets and slice configurations, PRISM accurately estimates slice thickness statistically indistinguishable from the true slice thickness, outperforming existing explicit kernel estimation approaches. Our contributions are:

1. A frequency-domain formulation of slice thickness estimation based on invariant spectral attenuation.
2. PRISM, a contrast-invariant, reference-based method that estimates slice thickness without alignment or paired data.
3. Extensive validation demonstrating accurate thickness estimation and improved downstream super-resolution performance.

2 Methods

The forward model Let $x(z)$ denote the continuous through-plane signal of the underlying anatomy, and let $h_\sigma(z)$ denote the continuous slice selection profile parameterized by σ. The acquired slice at location $z_m = ms$, where s is the slice spacing and m is the slice index, is given by

$$y[m] = \int_{-\infty}^{\infty} x(\zeta)\, h_\sigma(z_m - \zeta)\, d\zeta. \tag{1}$$

Thus, slice acquisition corresponds to evaluating the continuous convolution $(x * h_\sigma)(z)$ at discrete slice locations z_m. This model applies to oblique acquisitions as well, where the slice-selection axis defines the through-plane direction. In practice, PRISM is applied along the physical slice-selection direction obtained from the acquisition orientation metadata; because the comparison uses whole-volume spectra, no registration is required.

Following established MRI SR literature [7, 23], we model h_σ as a unit-area Gaussian profile. The FWHM of a Gaussian profile is defined in terms of its standard deviation: $\text{FWHM} = 2\sqrt{2\ln 2}\sigma$. The slice spacing s is assumed known from acquisition metadata, while σ is unknown.

We discretize the volume on a HR grid with spacing Δz and denote a single through-plane column of samples by $x[n]$. The same 1D slice-selection model is applied independently along the through-plane axis at every in-plane location, inducing a linear operator $y = A_\sigma x$ that acts on the full 3D volume; continuous convolution with h_σ is followed by sampling at slice locations, and slice spacing and thickness need not align with the discretization grid. Here and throughout, x, y, and $\tilde{x}$ written without an argument denote whole volumes, whereas $x(z)$ and $x[n]$ denote a single through-plane column in continuous and discrete form, respectively. For notational simplicity, we use A_σ to denote this discrete forward operator.

Spectral Signature of Slice Thickness Slice selection profiles act as 1D low-pass filters along the slice axis: increasing thickness produces steeper spectral decay. To characterize through-plane resolution, we define, for an arbitrary volume g, its 1D log power spectrum along the slice axis:

$$P[g](\omega) = \log\left(\frac{1}{HW}\sum_{i=1}^{H}\sum_{j=1}^{W}\left|\mathcal{F}_z(g)(i,j,\omega)\right|^2\right), \tag{2}$$

where H and W denote the in-plane voxel dimensions and $\mathcal{F}_z$ is the one-dimensional Fourier transform along the through-plane axis. The bracket notation $P[\,\cdot\,]$ emphasizes that the spectrum is a functional of whatever volume it is given; in PRISM it is evaluated not on raw intensities but on the rank-transformed volumes defined in the following section. Averaging over in-plane dimensions reduces sensitivity to anatomical variability, while the logarithm emphasizes differences in spectral decay. Matching this spectrum therefore corresponds to matching through-plane resolution.

Contrast-Invariant Structural Representation Reference and target images may differ in contrast, intensity scaling, and global brightness, so direct spectral comparison of raw intensities is not ideal. To suppress contrast-dependent variability while preserving structural information, we apply the rank transform [22] $R(\cdot)$ to both volumes prior to spectral analysis.

For each voxel location $\mathbf{u}$, the rank transform counts the number of neighboring voxels within a fixed $3 \times 3 \times 3$ neighborhood $\Omega(\mathbf{u})$ whose intensities are greater than the center voxel:

$$R(g)(\mathbf{u}) = \sum_{\mathbf{v} \in \Omega(\mathbf{u})} \mathbf{1}(g(\mathbf{v}) > g(\mathbf{u}))\,, \tag{3}$$

where $\mathbf{1}(\cdot)$ denotes the indicator function. Because the rank transform depends only on relative ordering, it is invariant to monotonic intensity transformations and robust to global contrast differences. However, blurring alters local ordering relationships, making the representation resolution-sensitive. This invariance is to monotonic, global intensity transformations rather than to every local contrast reversal; differing contrasts (e.g., T_1-w vs. T_2-w) can locally alter voxel ordering. Since PRISM compares 1D log power spectra averaged over the full in-plane field of view, such local changes perturb the baseline feature spectrum but are unlikely to reproduce the through-plane high-frequency attenuation induced by thicker slices, an assumption directly tested by our cross-subject, cross-contrast experiments.

Invariant Spectral Matching Let $\tilde{x}$ denote a HR reference volume. For a candidate slice thickness σ, we simulate acquisition using the discrete operator $\tilde{y}_\sigma = A_\sigma \tilde{x}$.

We compute the log power spectra of the rank-transformed volumes and define the spectral power error (SPE):

$$d(\sigma) = \sum_{\omega} \left(P[R(y)](\omega) - P[R(\tilde{y}_\sigma)](\omega)\right)^2 . \tag{4}$$

Equations 2 and 4 are computed over the whole volume, and co-registration is not explicitly required. However, we assume that $P[R(y)]$ and $P[R(\tilde{y}_\sigma)]$ have comparable anatomical statistics such that their spectra match when $A_\sigma \approx A$. Throughout this paper, we do not register any volumes to a template, but ensure common orientations among volumes.

The estimated Gaussian standard deviation is then obtained via

$$\sigma^* = \arg\min_{\sigma} d(\sigma). \tag{5}$$

Because the parameter space is 1D and well-behaved, we perform exhaustive search over a predefined range of clinically plausible slice thicknesses, which we choose as 20% to 120% of the slice separation. PRISM operates on whole-volume spectral statistics and requires no registration or voxel-wise correspondence between the acquired image and the reference. Each evaluation of $d(\sigma)$ requires a single vectorized matrix product, rank transform, and computation of a 1D

FFT along the slice axis. PRISM requires no iterative optimization, registration, or voxel-wise correspondence, and evaluates the 64 candidate thicknesses for a full 3D brain volume in approximately 10 seconds on a modern GPU.

3 Experiments and Results

3.1 Experimental setup

We evaluated PRISM on 50 T_2-w subjects from each of two publicly available brain MRI datasets: OASIS-3 [11] (1mm^3) and HCP-1200 [20] (0.7mm^3). All data are publicly available and de-identified; no new human-subject data were collected, and dataset use follows the respective OASIS-3 and HCP-1200 data-use agreements. A T_1-w image from a held-out subject in each dataset served as the reference volume. High-resolution (HR) isotropic volumes were treated as ground truth for simulation. Low-resolution (LR) acquisitions were simulated using the forward model in Section 2 for slice configurations $\tau \| \gamma \in \{3 \| 1, 4 \| 0, 5 \| 1.5, 10 \| 3.3\}$ mm, where τ denotes thickness and γ denotes the slice gap, and the slice gap is the difference between slice separation and thickness: $\gamma = s - \tau$. No in-plane degradation was applied, and the simulated thickness served as ground truth. HR resolution was modeled as Gaussian with FWHM equal to slice spacing. Target slice thicknesses were simulated using the sum-of-variances property of Gaussians to compute the additional blur required beyond intrinsic HR resolution.

We compared PRISM to the true simulated kernel (upper bound), a naive Gaussian assumption (thickness = spacing), and ESPRESO [9], a prior explicit kernel estimation method. Slice thickness estimation accuracy was measured as absolute error (mm) between estimated and true FWHM. For downstream super-resolution (SR), we report PSNR and SSIM between reconstructed volumes and ground truth HR images. Statistical significance was assessed using the paired Wilcoxon signed-rank test with Holm correction.

3.2 Slice thickness recovery

We first evaluated whether PRISM accurately recovers known slice thickness under simulation. Across all tested slice thickness–spacing configurations and both datasets, PRISM consistently produced thickness estimates close to the ground truth. Figure 2 shows the mean and standard deviation of the estimated kernels across subjects. The naive Gaussian approach overestimates thickness in the presence of slice gaps, while ESPRESO tends to underestimate thickness at larger slice spacings. In contrast, PRISM has low mean FWHM error to the true slice thickness across all tested conditions. Quantitatively, PRISM achieves lower mean absolute error than both naive and ESPRESO baselines in nearly all configurations. Performance remains stable across datasets, indicating robustness to anatomical variability. These results demonstrate that whole-volume spectral matching accurately captures the through-plane resolution induced by slice selection.

Table 1. Ablation results: mean ± std absolute errors in mm from the true slice thickness. The top-ranking (ϕ, d) pair within a resolution and dataset is **bolded**. Any (ϕ, d) pairs denoted by $*$ indicate no statistical significance ($\alpha = 0.05$) from the best pair according to the paired Wilcoxon signed-rank test with Holm correction. Acronyms: JSD: Jensen–Shannon Divergence. W1: Wasserstein-1. SPE: Spectral Power Error.

	ϕ	d	**3‖1**	**4‖0**	**5‖1.5**	**10‖3.3**
OASIS-3	Census Transform	JSD	0.15 ± 0.10	0.10 ± 0.07	$0.22 \pm 0.14^*$	1.31 ± 0.51
		W1	0.53 ± 0.29	0.53 ± 0.34	1.19 ± 0.79	4.49 ± 1.82
		SPE	0.12 ± 0.10	0.08 ± 0.07	$\mathbf{0.19 \pm 0.15}$	$0.31 \pm 0.24^*$
	Rank Transform	JSD	0.39 ± 0.32	0.42 ± 0.28	0.64 ± 0.47	2.29 ± 1.17
		W1	0.36 ± 0.31	0.38 ± 0.29	0.55 ± 0.41	2.20 ± 1.17
		SPE	$\mathbf{0.09 \pm 0.09}$	$\mathbf{0.06 \pm 0.06}$	$0.24 \pm 0.13^*$	$\mathbf{0.25 \pm 0.16}$
	Knutsson Edge Map	JSD	0.98 ± 0.63	1.27 ± 0.97	2.37 ± 0.83	3.99 ± 1.27
		W1	0.98 ± 0.62	1.26 ± 0.97	2.37 ± 0.83	4.08 ± 1.20
		SPE	1.90 ± 0.00	0.90 ± 0.00	2.86 ± 0.00	5.99 ± 0.00
	Fourier Phase	JSD	0.84 ± 0.49	1.59 ± 0.76	1.55 ± 0.74	4.04 ± 2.21
		W1	1.47 ± 0.42	2.19 ± 0.73	2.37 ± 0.48	4.69 ± 2.13
		SPE	0.95 ± 0.35	1.23 ± 0.95	1.90 ± 0.56	5.12 ± 1.88
HCP-1200	Census Transform	JSD	0.15 ± 0.22	$0.08 \pm 0.06^*$	$0.28 \pm 0.31^*$	$\mathbf{0.76 \pm 0.55}$
		W1	0.33 ± 0.34	0.28 ± 0.20	0.53 ± 0.33	1.00 ± 0.76
		SPE	$\mathbf{0.08 \pm 0.05}$	$0.06 \pm 0.04^*$	$\mathbf{0.16 \pm 0.12}$	$0.81 \pm 0.37^*$
	Rank Transform	JSD	0.30 ± 0.23	0.34 ± 0.26	0.49 ± 0.36	$0.86 \pm 0.69^*$
		W1	0.31 ± 0.25	0.33 ± 0.25	0.45 ± 0.33	$0.82 \pm 0.65^*$
		SPE	$0.08 \pm 0.05^*$	$\mathbf{0.05 \pm 0.04}$	$0.20 \pm 0.15^*$	$0.64 \pm 0.39^*$
	Knutsson Edge Map	JSD	1.33 ± 0.52	1.84 ± 1.01	1.33 ± 0.77	4.89 ± 1.40
		W1	1.37 ± 0.49	1.70 ± 1.10	1.33 ± 0.77	4.89 ± 1.34
		SPE	1.85 ± 0.00	0.85 ± 0.00	2.83 ± 0.00	5.98 ± 0.00
	Fourier Phase	JSD	0.77 ± 0.46	1.40 ± 0.80	1.60 ± 1.01	4.32 ± 2.14
		W1	1.31 ± 0.54	1.82 ± 0.78	2.07 ± 1.12	5.31 ± 1.64
		SPE	0.93 ± 0.38	1.49 ± 0.75	1.51 ± 1.08	4.92 ± 0.77

3.3 Ablation study

We investigated the impact of structural representation and spectral discrepancy metric on slice thickness estimation. We compared four structural transforms: the rank transform, the census transform [22], the Knutsson edge map [10, 18], and Fourier phase. We also evaluated three discrepancy measures: Spectral Power Error (SPE): mean squared error between log power spectra; Jensen–Shannon divergence (JSD) between feature histograms; and Wasserstein-1 (W1) distance between feature histograms.

Table 1 reports absolute slice thickness errors across configurations. Across both datasets and all tested resolutions, the combination of rank transform and SPE consistently achieves the lowest or statistically indistinguishable-from-lowest

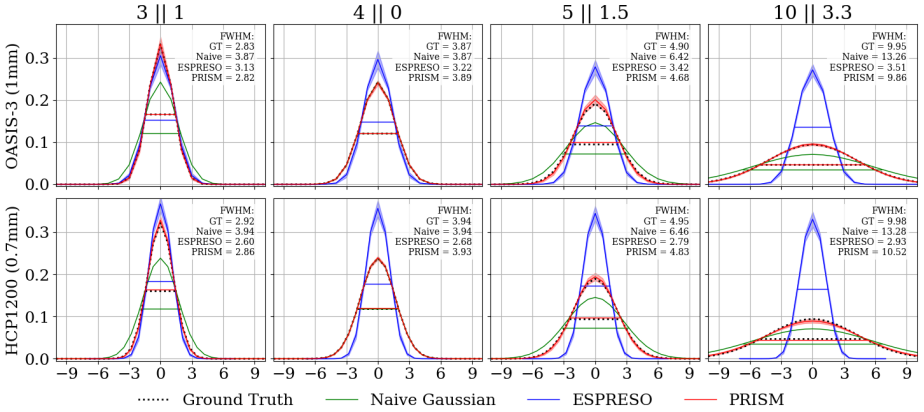

Fig. 2. Comparison between slice thickness estimation methods for the OASIS-3 dataset at 1mm (first row) and HCP-1200 dataset at 0.7mm (second row) for the slice resolutions 3‖1, 4‖0, 5‖1.5, and 10‖3.3. The mean slice profile for each method is drawn as a solid line with the standard deviation illustrated as a faded band around the profile means.

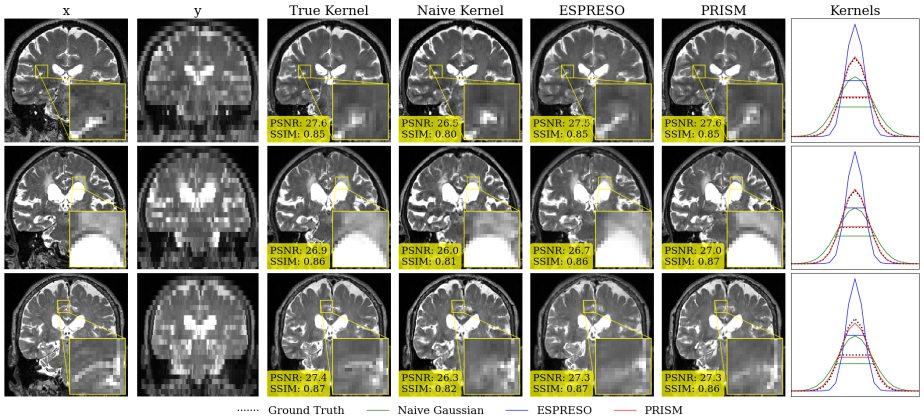

Fig. 3. Representative coronal slices from the downstream super-resolution experiment with corresponding estimated kernels. Zoomed insets show specific regions sensitive to the kernel.

error. Histogram-based metrics (JSD, W1) are generally less consistent across resolutions, and Fourier phase and edge-based features exhibit overall larger error. These findings support two key design choices of PRISM: (1) rank-based structural representation provides contrast robustness while remaining resolution-sensitive; and (2) direct spectral power matching most effectively captures slice thickness–dependent frequency attenuation.

Table 2. Downstream results: mean ± std PSNR and SSIM values for the four downstream approaches. An asterisk denotes statistical significance ($\alpha = 0.05$) between True and each method. An absent asterisk indicates no significant difference. Significance tests done with the paired Wilcoxon signed-rank test and Holm correction.

	True	Naive	ESPRESO	PRISM
PSNR	27.930 ± 0.763	$26.709 \pm 0.635^{*}$	$27.782 \pm 0.766^{*}$	27.905 ± 0.775
SSIM	0.870 ± 0.014	$0.814 \pm 0.012^{*}$	$0.874 \pm 0.017^{*}$	0.870 ± 0.018

3.4 Downstream super-resolution

We evaluated the practical impact of slice thickness estimation in a downstream super-resolution task. A pre-trained diffusion-based prior model [16] was used to implement the denoising diffusion null-space model [21] method for solving inverse problems. Model weights and seeds were fixed during SR; only the forward degradation model used during reconstruction differed. This experiment therefore isolates the effect of the estimated slice profile within a fixed reconstruction pipeline, rather than benchmarking super-resolution architectures or perceptual quality. We compared four configurations of the forward degradation model: the true slice profile, the naive Gaussian, ESPRESO, and PRISM.

Quantitative results on 50 OASIS-3 volumes are reported in Table 2. Using the naive Gaussian forward model results in significantly lower PSNR and SSIM. ESPRESO performance improved but remains slightly inferior to the true slice profile. PRISM achieved PSNR and SSIM statistically indistinguishable from the true kernel, while the naive Gaussian and ESPRESO are significantly different from the true kernel.

Qualitative examples are shown in Figure 3. Reconstructions using inaccurate kernels exhibit over-smoothing or residual aliasing, while PRISM-based reconstructions more closely resemble the ground truth HR anatomy. The ESPRESO result achieves comparable PSNR and SSIM to PRISM, but features near the edge of the brain and scalp have more hallucinations than PRISM when compared to the true kernel estimation and the ground truth x.

3.5 Real-World Evaluation

Finally, we evaluated PRISM on 28 real 2D MRI acquisitions from the OASIS-3 dataset for which slice thickness and spacing were known from metadata. This experiment tests behavior on real 2D acquisitions with known nominal thickness and gaps rather than validating the physical RF slice-selection waveform, which is unknown for these data. These data were $4\|1.2$, so slice gaps were present, but the exact shape of the slice profile was unknown. For each volume, we estimated slice thickness using the naive Gaussian assumption, ESPRESO, and PRISM. Thickness estimates were compared to acquisition-reported slice thickness.

PRISM estimates a mean slice thickness of 4.23 ± 0.05mm across the 28 subjects, compared to 3.35 ± 0.13mm for ESPRESO. The naive assumption (thickness

= spacing) yields 5.2 mm and therefore overestimates through-plane resolution in the presence of slice gaps. Consistent with simulation findings, the naive method overestimates slice thickness in the presence of slice gaps. ESPRESO consistently underestimates the thickness. PRISM produces thickness estimates closest to the reported acquisition parameters, but slightly overestimated. This is potentially because PRISM estimates the relative through-plane blur with respect to the reference image, not just the slice thickness, so any additional contributions to the through-plane point spread function will affect results. Although the exact continuous slice profile may deviate from an ideal Gaussian, these results indicate that invariant spectral matching robustly estimates slice thickness on real 2D-acquired volumes.

4 Discussion and Conclusion

We presented PRISM, a contrast-invariant method for slice thickness estimation in 2D multi-slice MRI based on whole-volume spectral matching to a reference image. By modeling slice acquisition as continuous convolution and leveraging the frequency-domain attenuation induced by through-plane blur, PRISM estimates slice thickness without requiring anatomical alignment, paired data, or iterative optimization. Across simulations, ablation studies, downstream super-resolution experiments, and real 2D-acquired volumes, PRISM consistently recovered slice thickness accurately and improved reconstruction quality relative to existing kernel estimation approaches. We model the slice profile as a Gaussian, used here as a compact one-parameter effective-FWHM approximation common in MRI rather than a claim that RF slice profiles are exactly Gaussian; the real-world results demonstrate that accurate slice thickness is estimated even when the true profile may deviate from this assumption. PRISM assumes that the through-plane spectral statistics are comparable between the low-resolution and reference volumes; the HR reference should therefore share comparable directional and anatomical spectral characteristics with the target. Because matching is performed on whole-volume spectra, no registration or voxel-wise correspondence is required. Conversely, severe motion, slice dropout, or interpolation artifacts in either volume may bias the spectral estimate. Future work includes investigating these sensitivities alongside extending PRISM to more flexible slice profile models and incorporating spatially adaptive spectral analysis to further improve robustness in the presence of extreme pathology or highly anisotropic anatomy.

Acknowledgments. This work received support from FG-2008-36966 (Dewey), DoD CDMRP HT9425-24-1-0785 (Dewey) and CDMRP W81XWH2010912 (Prince). Data used in this paper were sourced from open datasets; due to space constraints, please visit the citations of the appropriate data descriptor manuscripts for relevant information.

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

References

1. Altahawi, F., Pierce, J., Aslan, M., Li, X., Winalski, C.S., Subhas, N.: 3D MRI of the knee. In: Seminars in Musculoskeletal Radiology. vol. 25, pp. 455–467. Thieme Medical Publishers, Inc. (2021)
2. Bachrata, B., Bollmann, S., Grabner, G., Trattnig, S., Robinson, S.: Isotropic QSM in seconds using super-resolution 2D EPI imaging in 3 orthogonal planes. Proceedings 2022 ISMRM SMRT Annual Meet Exhibit Lond, London, England **2363** (2022)
3. Bell-Kligler, S., Shocher, A., Irani, M.: Blind Super-Resolution Kernel Estimation using an Internal-GAN. In: Advances in Neural Information Processing Systems. vol. 32 (2019)
4. Bernstein, M.A., King, K.F., Zhou, X.J.: Handbook of MRI pulse sequences. Elsevier (2004)
5. Del Grande, F., Hinterholzer, N., Nanz, D.: 3D MRI: technical considerations and practical integration. In: Seminars in Musculoskeletal Radiology. vol. 25, pp. 381–387. Thieme Medical Publishers, Inc. (2021)
6. Fergus, R., Singh, B., Hertzmann, A., Roweis, S.T., Freeman, W.T.: Removing Camera Shake from a Single Photograph. In: ACM SIGGRAPH. pp. 787–794 (2006)
7. Greenspan, H., Oz, G., Kiryati, N., Peled, S.: MRI inter-slice reconstruction using super-resolution. Magnetic Resonance Imaging **20**(5), 437–446 (2002)
8. Hacohen, Y., Shechtman, E., Lischinski, D.: Deblurring by example using dense correspondence. In: Proceedings of the IEEE international conference on computer vision. pp. 2384–2391 (2013)
9. Han, S., Remedios, S.W., Schär, M., Carass, A., Prince, J.L.: ESPRESO: An algorithm to estimate the slice profile of a single magnetic resonance image. Magnetic Resonance Imaging **98**, 155–163 (2023)
10. Knutsson, H.: Producing a Continuous and Distance Preserving 5-D Vector Representation of 3-D Orientation. In: Proceedings of the IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management. pp. 175–182. IEEE Computer Society (1985)
11. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., Raichle, M.E., Cruchaga, C., Marcus, D.: OASIS-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer Disease. medRxiv pp. 2019–12 (2019)
12. Lee, R., Li, R., Venieris, S., Hospedales, T., Huszár, F., Lane, N.D.: Meta-learned kernel for blind super-resolution kernel estimation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1496–1505 (2024)
13. Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Understanding Blind Deconvolution Algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence **33**(12), 2354–2367 (2011)
14. Pan, J., Hu, Z., Su, Z., Yang, M.H.: Deblurring Face Images with Exemplars. In: Computer Vision – ECCV 2014. pp. 47–62. Springer International Publishing, Cham (2014)
15. Pan, J., Ren, W., Hu, Z., Yang, M.H.: Learning to deblur images with exemplars. IEEE Transactions on Pattern Analysis and Machine Intelligence **41**(6), 1412–1425 (2018)
16. Remedios, S.W., Carass, A., Prince, J.L., Dewey, B.E.: Diffusion-Driven Generation of Minimally Preprocessed Brain MRI. arXiv preprint arXiv:2510.26834 (2025)

17. 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)
18. Stough, J.V., Glaister, J., Ye, C., Ying, S.H., Prince, J.L., Carass, A.: Automatic method for thalamus parcellation using multi-modal feature classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 169–176. Springer (2014)
19. Stringer, M.S., Lee, H., Huuskonen, M.T., MacIntosh, B.J., Brown, R., Montagne, A., Atwi, S., Ramirez, J., Jansen, M.A., Marshall, I., et al.: A review of translational magnetic resonance imaging in human and rodent experimental models of small vessel disease. Translational Stroke Research **12**(1), 15–30 (2021)
20. Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., et al.: The Human Connectome Project: A data acquisition perspective. Neuroimage **62**(4), 2222–2231 (2012)
21. Wang, Y., Yu, J., Zhang, J.: Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In: The Eleventh International Conference on Learning Representations (2023)
22. Zabih, R., Woodfill, J.: Non-parametric local transforms for computing visual correspondence. In: European Conference on Computer Vision. pp. 151–158. Springer Berlin Heidelberg (1994)
23. 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 (2021)
24. Zheng, H., Yong, H., Zhang, L.: Unfolded deep kernel estimation for blind image super-resolution. In: European Conference on Computer Vision. pp. 502–518. Springer (2022)